IEEE MICRO

DECEMBER 1987

GaAs Technology Matures



One-chip RISCs zoom to the forefront

THE COMPUTER SOCIETY OF THE IEEE

THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.

IEEE MICRO

Volume 7 Number 6 (ISSN 0272-1732) December 1987

Published by the Computer Society of the IEEE



On the Cover

Read more about the two gallium arsenide chips pictured on the cover by turning to pp. 3 and 8.

Cover illustration:
Jim Brummond

STAFF

Editor and Publisher
True Seaborn

Assistant Publisher
Douglas Combs

Managing Editor
Marie English

Assistant Editor
Christine Miller

Assistant to the Publisher
Pat Paulsen

Advertising Director
Dawn Peck

Art Director
Jay Simpson

Design and Production
Tricia Hayden

Membership Manager
Christina Champion

Circulation Manager
Paul Zive

Advertising Coordinators
Heidi Rex, Marian Tibayan

Reader Service
Marian Tibayan

Departments

From the Editor-in-Chief 2

About the Cover 3

MicroReview 4
Publishing technical papers

MicroNews 6
Neural networks interview; copyright hearing; trends in techniques, services

MicroLaw 86
Manufacturers' disclaimers of liability

MicroStandards 88
Future micro standards projects

New Products 92

Product Summary 100

Calendar 103

Advertiser/Product Index 104

Computer Society Information C3

Classified ads, p. 97; Change-of-Address form, p. 104; Reader Interest/Service/Subscription cards, p. 104A.

Feature Articles

A 32-Bit, 200-MHz GaAs RISC for High-Throughput Signal Processing Environments 8
Barbara A. Naused and *Barry K. Gilbert*
GaAs technology has matured sufficiently to allow fabrication of an entire RISC on one chip. GaAs also supports 200-MHz clock rates and 100-MIPS instruction rates.

A Distributed System for Real-Time Applications 21
Eli T. Fathi, Eloi Bosse, and *Jean Caseault*
This versatile network uses multiple microprocessors and a split-bus architecture to collect and process real-time data from a variety of sources and then transfer it to different destinations.

A General Heap Processor 29
Eduardo Sanchez, Patrick Sommer, Jacques Menu, and *Christian Iseli*
Switzerland's first 16-bit processor design may lead to a system capable of easing HLL implementation. It has already led to a better understanding of complex ICs.

A Fast Integer Binary Logarithm of Large Arguments 41
Reinhard Maenner
Simple and fast (10 μ), this new procedure requires no hardware and features a 10^{-6} approximation error.

Effective Implementation of a Parallel Language on a Multiprocessor 46
Thomas L. Sterling, Albert J. Musciano, Ellery Y. Chan, and *Douglas A. Thomae*
A new multiprocessor execution environment integrates semantics of parallel control with mechanisms for synchronization of concurrent tasks.

An Application-Specific Coprocessor for High-Speed Cellular Logic Operations 63
Robert E. Jenkins and *D. Gilbert Lee, Jr.*
Why develop a special coprocessor to perform one lengthy computation? Performance is much better than with equivalent software and fast-turnaround development times will soon be possible.

Software Design for Real-Time Multiprocessor VMEbus Systems 71
Walter S. Heath
Taming real-time systems can be a challenge. You'll gain flexibility and cut debugging time by adding additional processors and the right software.

Annual Index 81
Subject and author listing of 1987 *IEEE Micro* articles.

Circulation: IEEE MICRO (ISSN 0272-1732) is published bimonthly by the Computer Society of the IEEE: IEEE Headquarters, 345 East 47th St., New York, NY 10017; Computer Society West Coast Office, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720-2578. Annual subscription: $17 in addition to Computer Society or any other IEEE society member dues; $25 for members of other technical organizations. This journal is also available in microfiche form.

Postmaster: Send address changes and undelivered copies to IEEE MICRO, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720-2578. Second class postage is paid at New York, NY, and at additional mailing offices.

Copyright and reprint permissions: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limits of US Copyright Law for private use of patrons: those post-1977 articles that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 29 Congress St., Salem, MA 01970. Instructors are permitted to photocopy isolated articles for noncommercial classroom use without fee. For other copying, reprint, or republication permission, write to Reprints, IEEE MICRO, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720-2578. All rights reserved. Copyright © 1987 by the Institute of Electrical and Electronics Engineers, Inc.

Editorial: Unless otherwise stated, bylined articles and descriptions of products and services in New Products, Product Summary, MicroReview, MicroNews, and MicroLaw, reflect the author's or firm's opinion; inclusion in this publication does not necessarily constitute endorsement by the IEEE or the Computer Society of the IEEE. Send editorial correspondence to IEEE MICRO, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720-2578. All submissions are subject to editing for style, clarity, and space considerations. IEEE MICRO subscribes to The Computer Press Association's code of professional ethics.

BPA

IEEE Micro

Editor-in-Chief:
James J. Farrell III
VLSI Technology Incorporated*

Associate Editor-in-Chief:
Joe Hootman
University of North Dakota

Editorial Board
Shmuel Ben-Yaakov
Ben Gurion University of the Negev
Dante Del Corso
Politecnico di Torino, Italy
John Crawford
Intel Corporation
Stephen A. Dyer
Kansas State University
K.-E. Grosspietsch
GMD, Germany
David B. Gustavson
Stanford Linear Accelerator Center
Victor K.L. Huang
AT&T Information Systems
Barry W. Johnson
University of Virginia
David K. Kahaner
National Bureau of Standards
Jay Kamdar
National Semiconductor Corporation
G. Jack Lipovski
University of Texas
Kenneth Majithia
IBM Corporation
Richard Mateosian
Marlin H. Mickle
University of Pittsburgh
Varish Panigrahi
Digital Equipment Corporation
Ken Sakamura
University of Tokyo
Michael Smolin
Smolin & Associates
Richard H. Stern
Yoichi Yano
NEC Corporation

Magazine Advisory Committee
Michael Evangelist (chair),
Vishwani D. Agrawal, James J. Farrell III,
Ted Lewis, David Pessel, True Seaborn,
Bruce D. Shriver, John Staudhammer

Publications Board
J.T. Cain, (chair), Vishwani D. Agrawal,
J. Richard Burke, Gerald L. Engel,
Michael Evangelist, James J. Farrell III,
Lansing Hatfield, Ronald G. Hoelzeman,
Ted Lewis, Ming T. Liu, Ed Nahouraii,
David Pessel, C.V. Ramamoorthy,
Vincent D. Shen, Bruce D. Shriver,
John Staudhammer, Steven L. Tanimoto

*Submit six copies of all articles and special-issue proposals to James J. Farrell III, 8375 South River Parkway, Tempe, AZ 85284; (602) 752-6222; Compmail + j.farrell.

From the Editor-in-Chief

A very large number of comments appeared in the mailbag this month. As usual, I am grateful for your comments and suggestions. Our industry changes so quickly that communication from readers worldwide is extremely important in keeping *IEEE Micro* relevant. The reader response cards travel here from all continents except Antarctica.

The recent Computer Society magazine readership survey indicated that *IEEE Micro* had the most loyal readers (about 89 percent intend to renew this year). Albeit we won distinction in this category over the other CS magazines by a very small margin, this is still very gratifying. While we are vigorously looking for new subscribers, we are very attentive to the comments of those currently subscribing as well. When I see a well-written technical article submitted

The mailbag

"All of the articles are well written and simple to read...." R.A., Como, Italy

"I find it amazing that the June issue arrived two weeks BEFORE the April issue." H.D.A., Wellington, New Zealand (The June issue was mailed precisely two months after the April issue. I suspect that your April issue spent a good deal of time stuck in the bottom of a mailbag somewhere.—J.F.)

"I liked the TRON project—good stuff...BTRON is the way to go...Look out, DOS-TRON is coming!" J.N., Suva, Fiji Islands

"More articles on microprocessor applications for hydraulic/pneumatic circuitry." F.J.B.B., New London, CT (Being familiar with the New London area, I suspect that F.J.B.B.'s work involves nautical electronics.—J.F.)

"I liked the standards report." Bothell, WA

"I would like to see a bit more depth and detail in MicroLaw column." T.P.G., Torrance, CA (Richard Stern has written a detailed and extensive book on semiconductor copyright law, if you wish further study.—J.F.)

"I liked Clipper, Am29000, Micro-Review, and MicroLaw. I would like more tutorials on microprocessors and VLSI." K.W., Nashville, TN

"Compare the 80387 to the Weitek 80387 replacement....(*Micro*) New Products Department is a wimpy substitute of a concept. For the real thing, check out (several commercial trade publications listed)." R.G., Randolph, NJ (Most comments indicate that our new products are right on target. We get the same news releases everyone else gets. Being bimonthly, our major problem in this department is timeliness.—J.F.)

"I liked little—collection of pseudo-intellectuals. I disliked extreme theory—wrong magazine. I would like to see clean applications articles by American authors." (origin withheld) (*IEEE Micro* reviews all material for publication without regard for the sex, race, religion, handicap, or national origin of the author.—J.F.)

"I liked New Products section, IEEE standards news." M.B., Brisbane, Australia

"I liked the new generation of microprocessors, introduction to the Clipper architecture, the 80387 and its applications." W., Vernon, France

"'System Considerations in the Design of the Am29000' was superb. I would like to see more articles such as the Am29000. How about articles on MIPS-X, Acorn's RISC processor, etc.?" V.L.D., Bangalore, India (Your suggestion is well taken. We have scheduled a special issue on this for June 1989, but perhaps we should address some RISC processor issues sooner. – J.F.)

"I liked the lack of HYPE in the microprocessor articles; reveals the honest

from Saudi Arabia receive complimentary comments from such geographically diverse locations as Mexico, the United Kingdom, and the Fiji Islands, I know that *IEEE Micro* is serving its readers.

Final note: It is with great sadness that we note the death of W.H. Brattain, one of the inventors of the transistor and a giant of science. He will be missed and remembered.

Best regards,



Jim Farrell

excitement of this field. THANK YOU. I would like to see monthly issues, please." R.J., Louisville, KY (I couldn't agree more; unfortunately, monthly issues are not presently in our plans.—J.F.)

"I would like to see more on VLSI and algorithms for signal processing...." S.G., Newcastle upon Tyne, UK

"I liked 'A Synthetic Instruction Mix.' It was a very good article that helps in choosing a microprocessor...more articles on mass use (e.g., vending) machines." M.J.H., Vina DelMar, Chile

"I liked the speech analysis article." J.N., Suva, Fiji Islands

"I liked the article on speech analysis and synthesis." R.H.R., Manchester, UK

"All of the articles in this (June) issue were excellent." J.G.P., Montreal, Canada

"The article on speech synthesis is very good." A.R.B., Mexico City

"I liked 'The Architecture of a Capability-Based Microprocessor System'...more on data communications systems." B.G.Y., Inchon, Korea

"I liked 'The Architecture of a Capability-Based Microprocessor System.'" G.P.M., Warsaw, Poland

"I liked 'The Architecture of a Capability-Based Microprocessor System.'" A.R.V., Bruyeres-le-Chatel, France

"I disliked the articles of a research-oriented nature. Let's see more applications." R.T., Newark, DE (More applications-oriented articles are planned.—J.F.)

"I need to obtain more information on products (catalog, price, etc.)." T.V.N., Los Alamitos, CA (Unlike some commercial trade magazines, contributed technical articles are not followed by a "bingo" number for more information. Please contact the author directly.—J.F.)

"Dr. El-Imam's paper has an error on page 8. (D/A converter should be A/D converter)...." C.L.B., Arlington, VA (You are correct. Several other readers noted the error also.—J.F.)

"I liked MicroLaw...." R.P.B., Phoenix, AZ

"Please print name AND date on the bottom line together on EACH page, e.g., IEEE MICRO/August 1987 – Be proud of your work! Show it!" R.L.Z, Munich, West Germany (Excellent suggestion. I will pass this on to the Publications Board and the Publisher with my full endorsement.—J.F.)

"I liked all the articles." A.P, La Plata, Argentina

"I would like to see more articles of the quality of Yousif A. El-Imam. This one was excellent as it informs but also teaches." A.R.B, Mexico City (I have received several very favorable comments about this article. Author El-Imam did an excellent job.)

"I liked elegant design in capabilities article; MicroLaw desperately important." D.E., Perth, Australia (To the best of my knowledge, the MicroLaw-type legal column is unique to *IEEE Micro*.—J.F.)

More about the cover

Complete systems fabricated in gallium arsenide are possible in the very near future. Such systems can be very valuable because they function well in extremely hot or cold environments.

Researchers involved in one digital signal processing project found that the fabrication of a GaAs microprocessor would ease the integration of complete systems. After considering a number of microprocessor architectures, they selected the RISC, or reduced instruction set computer, as the best approach. This research, funded by the US Defense Advanced Research Projects Agency known as DARPA, ultimately produced single-chip designs based on two different GaAs technologies.

On the right of our cover you can see a microphotograph of a 32-bit Data Path demonstration chip, detailing 76 percent of the complete CPU chip. Our thanks to McDonnell Douglas Astronautics Company for supplying this picture.

The microphotograph appearing on the left side of our cover depicts a complete prototype CPU chip from Texas Instruments Incorporated. *IEEE Micro* also extends its thanks to TI and to the Mayo Foundation authors who prepared the article, "A 32-Bit, 200-MHz GaAs RISC for High-Throughput Signal Processing Environments." This article, which appears on page 8 of this issue, relates the background, current status, and fabrication expectations of the GaAs projects.

MicroReview

Richard Mateosian
2919 Forest Avenue
Berkeley, CA 94705
(415) 540-7745

Producing technical papers

Publication is the medium by which we communicate the results of our technical work to others. We do this to advance our chosen fields and to obtain the recognition of our peers. Sometimes this publication is an article in a professional journal like *IEEE Micro*; sometimes it's an internal document for a company or governmental agency; sometimes it's a set of lecture notes. Technical publication takes many forms, but they all have this in common: The greater the clarity and the better the appearance of the document, the greater the chance of its achieving the desired result.

A journal like *IEEE Micro* takes much of the responsibility for the appearance of articles appearing in it, but for conference proceedings and other similar publications, authors are often responsible for providing "camera-ready" copy, which is simply reproduced unchanged in the final book. Furthermore, in this era of desktop publishing, standards have risen for all forms of communication. Readers have come to expect reasonably well-designed and clearly printed memos, notes, manuals, and so forth. They are not as willing as they once were to plow through scrawled notes or poorly typed and badly reproduced papers. If you want to build and keep an audience for your work, you can't just let the content speak for itself; you must turn your attention to the quality of the presentation.

Of course, the content of your work is still the foundation of your technical reputation. Thus, the ideal tool for producing your work will allow you to achieve a high quality of presentation without forcing you to become a "wizard" for the tool itself. This month I've looked at two programs for the Macintosh that are capable of handling the kinds of papers that *IEEE Micro* readers are most likely to be called upon to produce. Although both of these tools need to be learned and provide ample opportunity for wizardry, they are capable of producing high-quality output for the casual user. One uses the popular "What You See Is What You Get" (WYSIWYG) approach, while the other is a formatter, i.e., it processes input consisting of text interspersed with formatting commands.

The trade-offs between these two approaches are well known. WYSIWYG is easy to use and fast, especially for small jobs, but what you see on your screen may not be what you get on your laser printer or typesetting machine. Formatting commands give precise control, but you don't know what you're going to get until you see the printed output. Also, formatting systems tend to be slower than WYSIWYG systems.

The following discussion is based upon my use of these products on a Macintosh Plus with a Data Frame XP20 hard disk, interfaced through the SCSI port.

Word 3.01, Microsoft (Redmond, Wash., 1987, $395.00)

Word 3.01 is a pretty impressive WYSIWYG word processor. It has many features, but you have to hunt for them. Word 3.01 comes with a 150-page tutorial book, which only hits the highlights. After that you have to dig everything out of the 458-page manual, which is arranged like an encyclopedia—alphabetically by topic with no logical structure and no introductory or tutorial material. The on-line help facility is extensive and convenient, but you have to know what you're trying to do before you can use it.

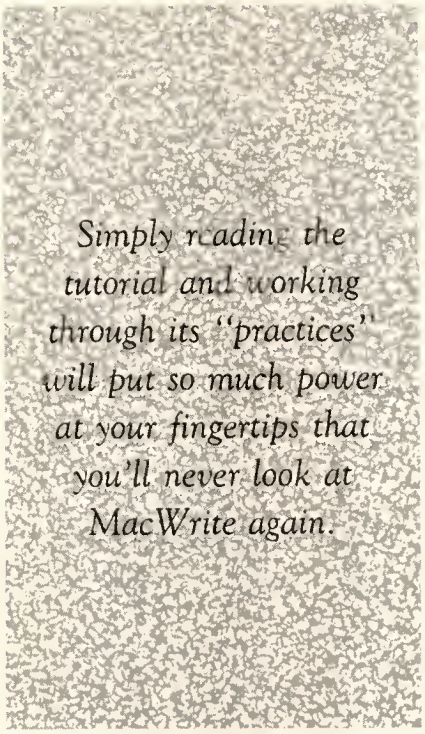
Simply reading the tutorial book and working through the "practices" in it will put so much power at your fingertips that you'll never look at MacWrite again. Then you can browse in the manual at your leisure or turn to it when a specific problem needs to be solved. That's the stage I've reached, and I've been pleased by the readability of the

manual and the ease of finding an appropriate feature for any problem I've had to solve. The only thing I've looked for that I haven't found is a macro facility. I didn't have a specific problem that I was trying to solve with macros, so I can't say whether an alternative was available.

Word lets you do just about anything that you need to do to produce a paper. You can start by outlining the material and then transform the outline into a paper, always keeping the outline structure available as a hidden part of the document. You can assign names to "styles," which are collections of paragraph and character formats, thus achieving uniformity of appearance within your paper or between papers. Footnoting, subscripting, superscripting, creating page headers and footers, page numbering, date and time stamping—all are easily accomplished with Word. You can format tabular material easily and even perform simple computations on selections of table entries. Glossaries allow named blocks of text and graphics to be placed in your paper with a few quick actions. A standard dictionary, which you can augment with your own special terms, supports spelling checking. You can prepare graphics using any Macintosh drawing or painting program (or charts using Excel), switch between these programs and Word with Apple's Switcher (supplied with Word), and place these exhibits in your paper.

You can easily resize and manipulate graphics with Word. If you're a Postscript expert, you can even include Postscript commands in your paper, allowing you to control the laser printer directly.

One area in which Word deviates from pure WYSIWYG and adopts the formatter approach is the production of formulas. You can type simple subscripts and superscripts, sums and products, and equations directly, but Word employs a functional notation for more complex forms (e.g., fractions, integrals, indexed sums and products, roots, matrices). Word has a general facility for toggling between two screen display modes: the WYSIWYG display and a display that shows visible representations of spaces, tabs, carriage returns, and other such characters. The functional notation is entered in the latter of these modes, and toggling to the former causes the specified formula to be displayed in final form. Nothing could be more convenient.



I've focused here on features relevant to producing technical papers and simple manuals. I haven't mentioned multicolumn printing, form letters, automatic hyphenation, page previewing, sorting, searching, interchange formats, and much more. There is even an entire set of keyboard commands that allows "mouse-free" operation for those already familiar with Microsoft Word for the MS-DOS environment. Microsoft's stated goal of keeping its MS-DOS and Macintosh products "in sync" is another plus for Word.

This unrestrained praise of Word should not be taken as a recommendation of Word over any competing product. This is not a comparison study. Other word processors may provide similar functionality or a better price/performance ratio. Simply take this as a baseline: Given what Word offers, there is no need to put up with less.

TeXtures, Addison-Wesley, (Reading, Mass., 1987, $495.00)

We have to begin with a discussion of etymology and pronunciation. TeXtures is spelled with an uppercase X in the middle because its name is derived from Donald Knuth's TeX. TeX is an approximation used on ordinary keyboards for TEX. This is the uppercase form of the Greek word spelled *tau, epsilon, chi,* which means art and technology. Thus, says Knuth, TeX is pronounced to rhyme with *blecchhh*, and all words formed from TeX share that pronunciation. One who masters TeX is a TeXnician, not a TeXpert, and the name of the product being described here is pronounced *tecchhhtures*, not textures. TeXtures is named as a contraction of "TeX plus pictures," because it marries the powerful formatting of TeX with the graphics capabilities of a Macintosh.

TeX is an edifice far too grand to be described adequately here, but I'll try to talk about some of its simpler features. TeXtures is an implementation of TeX for the Macintosh.

Every implementation of TeX has as its core the formatter described in *The TeXbook*, the first of Knuth's five-volume set called *Computers and Typesetting*. Thus, *The TeXbook*, a work of 483 pages, is included with TeXtures and forms the major part of the documentation. The remainder of the documentation is the 123-page user's guide, which describes the part of the package that is specific to this implementation of TeX. Thus, the user's guide deals with a small built-in editor, the facility for handling graphics, and the mechanism for previewing the final output on the Macintosh screen.

The editor is primitive, but adequate. Its purpose is rapid correction of errors uncovered during the formatting process. The file of text and formatting commands supplied to TeX is normally created with a standard editor, then imported by TeX.

Graphics are handled with the "picturebook," which works like a Macintosh scrapbook. You can copy graphics from other Macintosh application programs and then transfer them from the clipboard to the picturebook, where they can no longer be manipulated. Graphics are stored by name in the picturebook, along with precise size information for placing the graphic in the document. You can easily specify scaling to any arbitrary proportion of actual size.

You can preview the final output on the Macintosh screen one page at a time at any magnification from 10 percent to 500 percent of actual size. A magnifi-

Continued on p. 89

MicroNews

MicroNews features information of interest to professionals in the microcomputer/microprocessor industry. Send information for inclusion in MicroNews ***one month before cover date to Managing Editor,*** IEEE Micro, *10662 Los Vaqueros Circle, Los Alamitos, CA 90720-2578.*



"We don't need hype," Robert Hecht-Nielsen

Neurocomputing: a new information processing paradigm

Christine Miller, Assistant Editor

Whaaat? They're building a brain? A machine that sees, talks, and thinks? Are we ready for this?

IEEE Micro wondered and decided to approach Robert Hecht-Nielsen of the Hecht-Nielsen Neurocomputer Corp. We wanted to find out exactly what's happening in neural network technology and what all this means to the microcomputer industry.

Neural networks model the way the brain encodes and processes information. They recognize human faces, teach themselves to read and speak, learn from experience, and perform a variety of other pattern-recognition tasks that have baffled conventional computers.

Hecht-Nielsen is one of the first of many to develop neural networks for commercial applications and has built a neurocomputer. We plan to present other approaches in future issues.

Observers have called neural networks an emerging field. Do you agree?

Yes. Neural networks is one of the fastest growing areas of information processing, if not *the* fastest growing area. As more people see the capability demonstrations, they become aware that neurocomputing is a totally new information processing paradigm with numerous applications.

Since neurocomputers are not programmed, there is no need for software or software development. Development time and cost is lower, and networks can be implemented in simple hardware, with parallel hardware. All in all, implementation cost of neural network applications is much lower than costs for developing other applications.

Another reason neural networks are growing is that they are easy to learn for almost anyone already in information processing. I predict that neural networks will become the approach of choice for use in problem areas where neurocomputing is the application.

Currently, neurocomputing is a weak technology because we only know how to apply it to a rather narrow range of problems. However, this range is expected to broaden as we learn more.

To your knowledge, how many companies are presently developing and/or marketing commercial applications for neural networks?

I'm aware of at least 60 or 70 companies developing applications. Industries involved in this research are defense, financial, entertainment, telecommunications, electronics, and aerospace, to name some.

DARPA [the Department of Defense Advanced Research Projects Agency] through Dr. Ira Skurnick is very active in researching radar, sonar, weapons, and novel sensing applications.

The financial industry is using neural network experimentation with high-dimensional analysis of data. For instance, they are using neural networks to assess problems with credit line usage, or trouble signs on loan applications, or to score credit applications.

Neural networks can often be applied where there is an enormous database to be analyzed, like those involved in forecasting the success of a new movie or a new product. Other areas of application include special effects in films, automotive manufacturing, airline fare management, securities, and robotics.

What does all this mean to the microcomputer industry?

Microcomputers are the ideal host platforms for neurocomputing and will be used with neural networks because of their flexibility. For instance, IBM PC ATs and compatibles can handle a wide variety of hardware and software attachments and can help provide numerous applications in a cost-effective manner.

Microprocessors are the universal path for applying neural networks across a broad base. Other ways are prohibitive in cost.

What approach are you using to implement neural networks?

The neurocomputer is a coprocessor board that plugs into the host PC, which controls the neurocomputer. The host PC is loaded with a User Interface Subroutine Library. Using the UISL, users can call neural networks from software running on the host microcomputer as if they were subroutines. This makes neural networks easy to use anywhere they might be appropriate in a particular information processing application. This coprocessor approach allows users to freely mix programmed computing and neurocomputing so that each can carry out the processing it does best. Actually, it's simple to run.

Also involved is Axon, the machine-independent language for describing neural networks, existing or new. Just as an algorithm can be expressed in software, neural networks can be described in neurosoftware languages such as Axon. This approach allows the design of a combined programmed computing/neurocomputing system to be documented and maintained in a manner similar to current software maintenance procedures.

We have created network packages—most customers are only interested in four or five main networks—to be loaded into the neurocomputer on disk. The package usually provides only a general description of the network; users can tailor it to fit their needs.

On what theories are your implementations based?

Our Anza product is not specifically based on any one of the theories, but will work with all neural networks: Rumelhart backpropagation, Hopfield, and Grossberg counterpropagation, and other networks by Grossberg and Kohonen. One package, for instance, uses the Spatiotemporal Pattern Recognition Network by Grossberg.

[*Rumelhart, Hopfield, Grossberg, and Kohonen, among others, pioneered the development of neural networks.—Ed.*]

What is your source of funding?

We have had two rounds of financing: (a) the seed round with some 15 private investors; and (b) syndicate venture capital company investment. I am not at liberty to divulge the amounts involved.

"It is irresponsible and extremely grandiose to think that anything can function like a brain. That is a long way off."

More about applications. How are they selected?

Our company offers courses to train individuals in neural network capabilities and limitations. These "domain experts" then go out into their areas of expertise, whether it be commercial or governmental, and identify possible applications. We will also assist in this process.

Speaking of domain experts, is there a potential for use with artificial intelligence and/or expert systems?

A great potential. These are very compatible technologies; they don't compete. AI people are applying neural computing alongside knowledge engineering technology. I think this area will be very fruitful over the next two years or so.

A group known as the "Connectionists" believe that computers will act like brains when they are built like brains. What is your comment?

Connectionists are a part of the neural network community who are trying to make a point: If you want behaviors more like humans or animals, you have to use a different paradigm than before. But it is irresponsible and extremely grandiose to think that anything can function like a brain. That is a long way off.

We're not talking about brain capabilities. AI has had its hype—machines that can see, hear, whatever, have been promised, but never delivered. Our present capabilities are modest in comparison, but still impressive. We don't need hype.

Will the new 32-bit operating systems affect the development of neural networks?

We're excited. Hecht-Nielsen will be fully compatible. The software on a microprocessor is limited now in terms of memory and pointer limitation. The universal applicability of microcomputers will be assured by unlimited address space and other features.

What else is new on the horizon for you?

We are now beta-testing a new package called the Adaptive Resonance Network for hypothesis testing developed by Grossberg and Carpenter. As an illustration of how it works, consider an Automatic Teller Machine that can speak and listen to a client. The problem has been getting the software to focus on one character at a time. With this package, the teller machine focuses on the voice of the person as distinct from background noise, or in sonar it can focus on just one vehicle. We expect this package to be available the first quarter of next year.

IEEE Expert Call for Papers

IEEE Expert invites articles on **AI Applications** including **MIS/Financial Systems, Tools and Techniques, Health Care, Engineering/Manufacturing,** and **Education**—on **Technology** including **Real-Time Systems, Databases, Vision/Robotics, Software Engineering, Search, Natural Languages,** and **Knowledge Engineering**—plus **Book Reviews, Conference Coverage, Short Subjects,** and papers on **PCs, Products,** and **Resources.** Please submit articles to Editor-in-Chief David Pessel, BP America, 4440 Warrensville Center Rd., Cleveland, OH 44128. Please submit other papers to Henry Ayling, Managing Editor *IEEE Expert,* 10662 Los Vaqueros Circle, Los Alamitos, CA 90720. Book reviewers please contact Associate Editor K.S. Shankar, Federal Systems Division, IBM Corporation, 3700 Bay Area Boulevard, Houston, TX 77058.

Continued on p. 90

FEATURE

A 32-Bit, 200-MHz GaAs RISC

for High-Throughput Signal Processing Environments

Barbara A. Naused and Barry K. Gilbert
Mayo Foundation

GaAs technology has matured sufficiently to allow fabrication of an entire RISC on one chip. GaAs also supports 200-MHz clock rates and 100-MIPS instruction rates.

The first complex single-component microprocessor fabricated in gallium arsenide (GaAs) is now being developed as a major core element in a project known as the Advanced Onboard Signal Processor. Funded jointly by the Defense Advanced Research Projects Agency (DARPA) and the US Air Force, the AOSP will be assembled by 1990 entirely with GaAs components as a demonstration that digital GaAs technology will be sufficiently mature to be used in signal processing systems.

In the early 1980's, DARPA asked the Special Purpose Processor Development Group of the Mayo Foundation to identify a microprocessor architecture suitable for implementation as a custom-designed monolithic GaAs IC by the late 1980's. A number of commercial, aerospace, and Department of Defense microprocessor architectures were reviewed for their suitability. Here, we compare their strengths and deficiencies in the AOSP application context and discuss the selection of a RISC architecture for GaAs implementation.

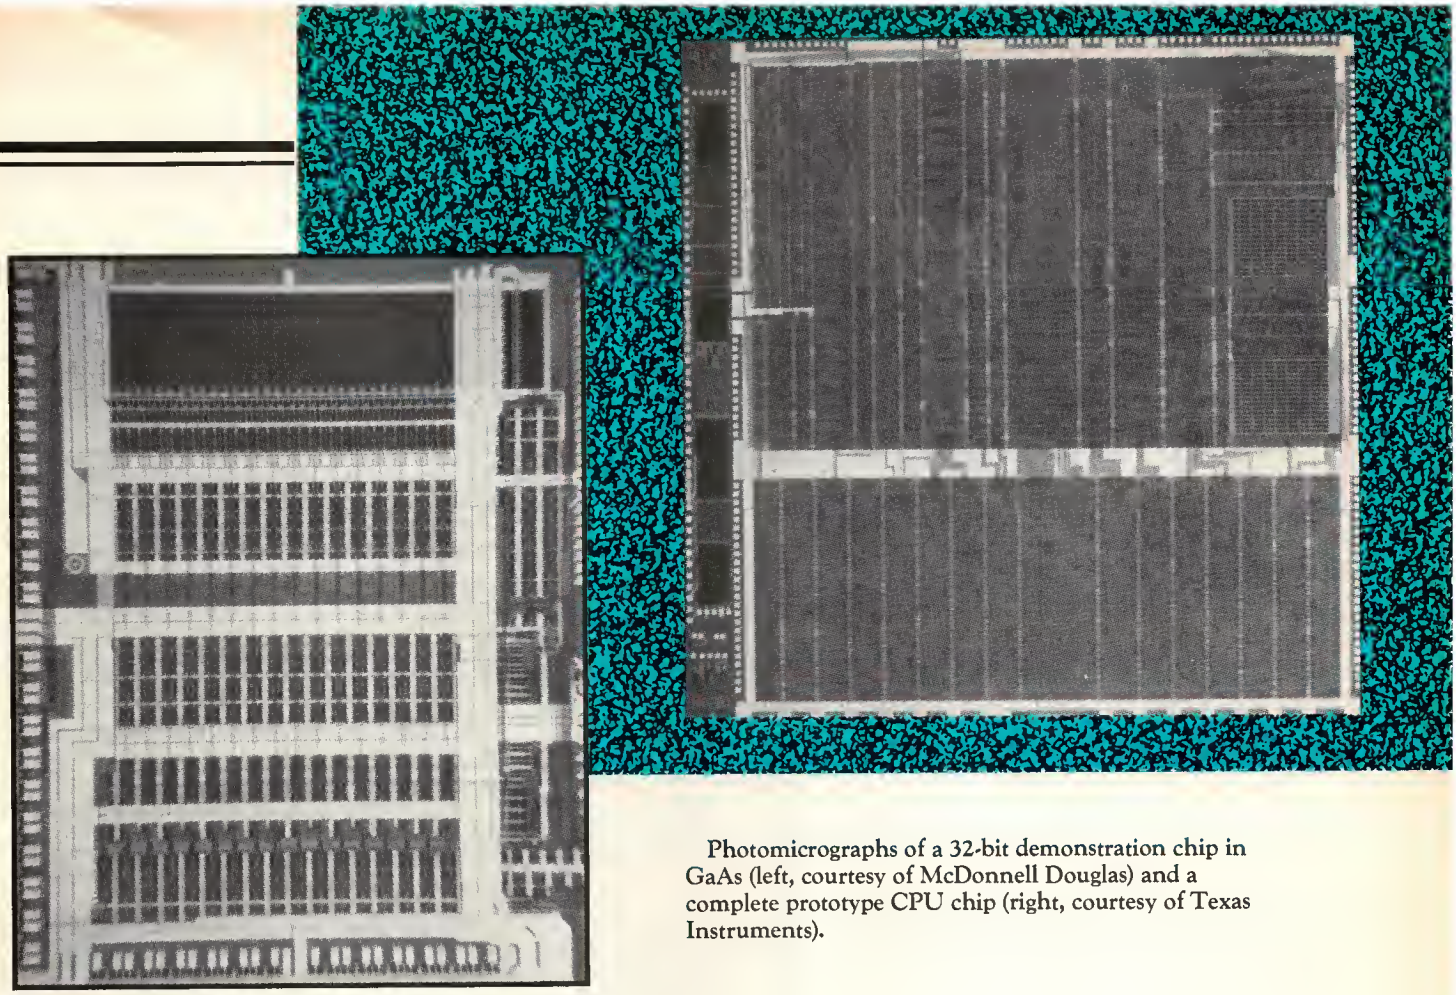
The AOSP

The AOSP is a general-purpose, distributed signal processor presently under development in silicon technology by Raytheon Corporation for use in several DoD signal processing applications. The architecture of the AOSP is based on a network of processing elements, called array computing elements (ACEs), which can all be identical, or can be a variety of types to perform several different functions. A distributed operating system controls the network, which is specialized for signal processing. The ACEs communicate through an efficient interconnection network which also allows spare ACEs to be substituted for faulty ones in real time (Figure 1). Each ACE contains two sections:

- *the Network Control Unit* (NCU), which is responsible for communication with the network and control of the ACE, and
- *the Application Processing Unit* (APU), which is optimized for execution of signal processing applications and also interfaces directly with the outside world.

Each NCU consists of an interelement bus interface and a control processor. The APU contains a

0272-1732/87/1000-0008$01.00 © 1987 IEEE



Photomicrographs of a 32-bit demonstration chip in GaAs (left, courtesy of McDonnell Douglas) and a complete prototype CPU chip (right, courtesy of Texas Instruments).

signal processor, referred to as the Macro Function Signal Processor (MFSP), a system I/O unit, and a data processor for scalar operations. A Motorola 68010 microprocessor functions as the control processor in the NCU and the data processor in the APU, as illustrated in Figure 2.

As the design of the AOSP progressed in silicon, DARPA management realized that certain applications would require very high clock rates, extreme radiation hardening, and/or very low power dissipation and that a GaAs implementation of the AOSP might achieve these combined features.[1]

Several studies conducted during 1981 and 1982 assessed the level of technology required to achieve a radiation-hardened, low-power signal processor by the end of the decade. These studies used the AOSP architecture as a baseline. Results indicated that GaAs IC technology would have to be sufficiently sophisticated to permit fabrication of configurable cell or gate arrays in the size range of 4000 to 6000 equivalent gates, as well as static RAMs (SRAMs) of at least 16K bits.

Assessments of the AOSP architecture indicated that—with one exception—an entire AOSP ACE could be fabricated using only the gate arrays and SRAMs. To assure that these technologies would be in place by the mid-1980's, DARPA embarked upon an ambitious plan to fund the creation of three pilot fabrication line facilities for GaAs chips. The products of the first two pilot lines were to be gate arrays of at least 6000-equivalent-gate complexity, and SRAMs of 16K-bit complexity. These projects are doing well. A Rockwell/Honeywell team demonstrated 5500 gate arrays in late 1985, and both Rockwell and McDonnell Douglas demonstrated 4K-bit SRAMs in early 1986. Rockwell demonstrated fully functional 7000 gate arrays in mid-1987 and fully functional 16K-bit SRAMs in September 1987.

A single exception to the assembly of an entire AOSP ACE based upon GaAs gate arrays and SRAMs remained: the requirement for a microprocessor to serve as the control and data processor elements of the AOSP ACE, as we earlier discussed. In order to achieve the full benefits of digital GaAs technology, the entire AOSP ACE had to be fabricated with GaAs components, including the microprocessors. However, further studies indicated that no modern microprocessor could be assembled efficiently with gate arrays. A custom IC was necessary to achieve optimum performance. Because GaAs IC technology was then, and remains, relatively immature in comparison to silicon technology, the device densities of GaAs gate arrays and custom components are considerably less than in silicon. As a result, a major constraint was placed on the design of a GaAs microprocessor. It would have to require less than 10,000 equivalent gates to allow its placement on a single custom chip.

The AOSP application placed additional constraints on the design of a microprocessor:

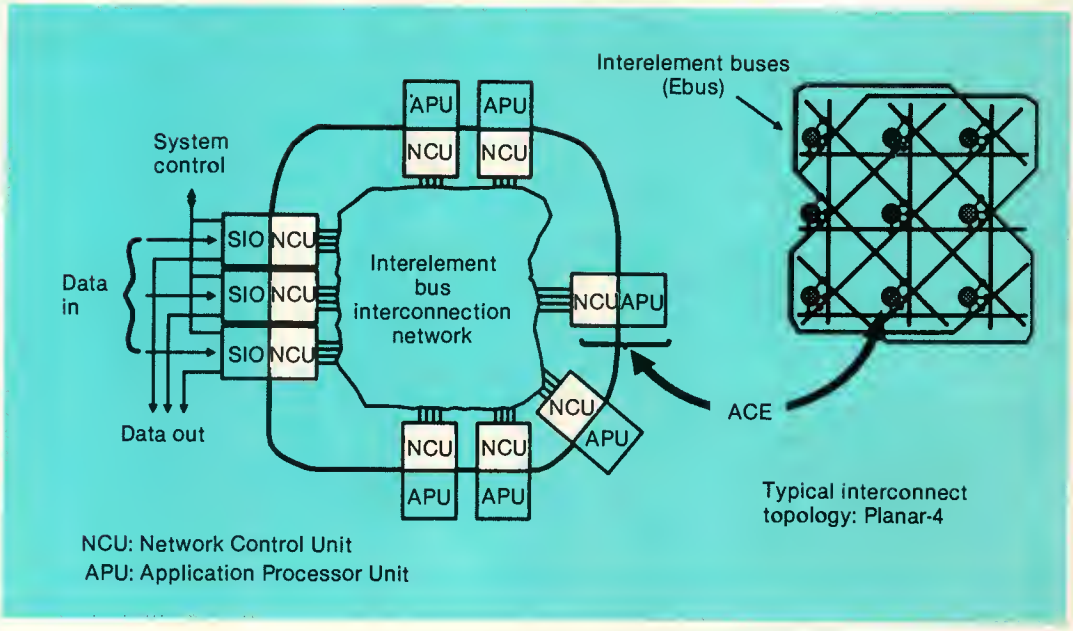


Figure 1. The AOSP system architecture developed by Raytheon Corp.

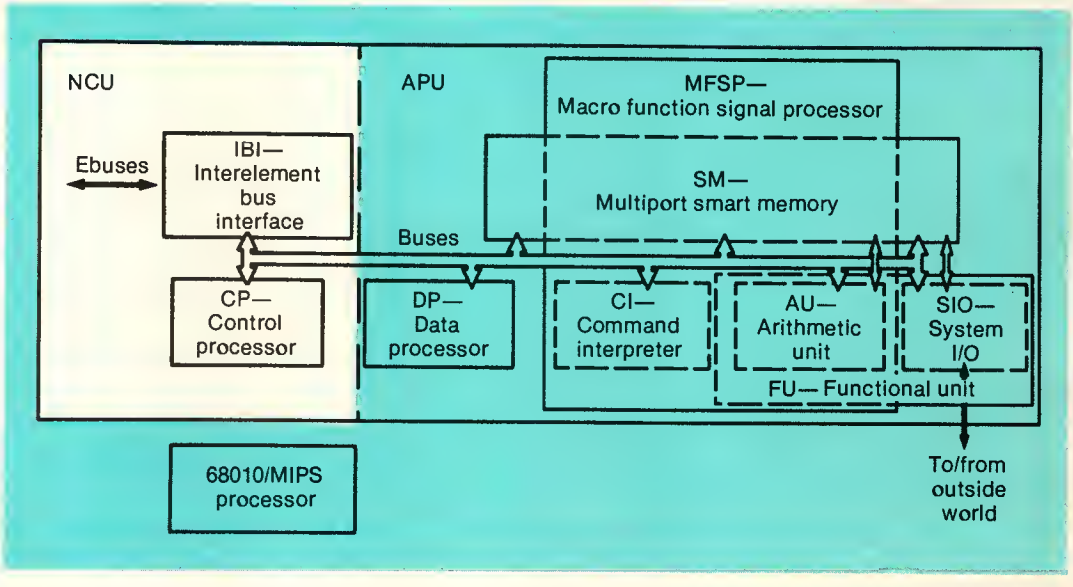


Figure 2. AOSP Array Computing Element (ACE) showing both the Network Control Unit (NCU) and the Application Processor Unit (APU).

- a 32-bit architecture;
- a 24- to 32-bit address bus to support at least 16M bytes of direct memory access (DMA);
- an efficient I/O handler;
- an efficient interrupt handler;
- a high-instruction execution rate; and
- the capability of floating-point operations in hardware.

If necessary, the floating-point hardware could be a separate coprocessor chip. Accordingly, DARPA tasked us to identify a 32-bit microprocessor architecture for full implementation in GaAs in time to achieve an all-GaAs AOSP brassboard demonstration in 1990.

Microprocessor architectures

Although a microprocessor could have been designed specifically for GaAs implementation, we did not consider this type of approach attractive without guaranteed software support. The possible use of an existing architecture therefore added another constraint to the GaAs microprocessor selection process. A microprocessor design that was nonproprietary, readily available, and well-documented was required. These constraints formed the basis of our study.

Commercial architectures. We first investigated commercially available microprocessor architec-

Table 1.
Possible architectures for use in a GaAs microprocessor.

	Gate count < 10K	32-bit architecture	Non-proprietary	DMA ≥ 16M bytes	Execution > 1 MIPS
Standard architectures					
Mil-Std 1750A (TI)			●		●
TI 9900					
MC68010 (12 MHz)		●		●	●
NS16032		●		●	
Nebula		●	●		
Intel iAPX432		●		●	
Rockwell AAMP				●	
RISC architectures					
RISC II - Berkeley		●	●	●	
IBM 801	?	●		●	●
Inmos transputer		●			●
MIPS - Stanford	●	●	●	●	●

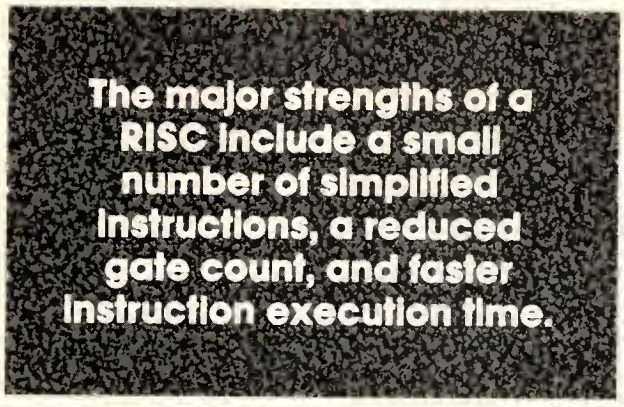
tures—in particular, the Motorola 68010. Because this microprocessor is used in the silicon version of the AOSP under construction by Raytheon, it would have provided a straightforward target architecture. However, the gate count of this processor is too large for a single-chip implementation in GaAs at projected late-1980's integration levels. Further, the design for the 68010 is proprietary and would not be made available. Raytheon discovered additional drawbacks of the 68010. Wait states are incorporated in the initialization of I/O operations, and the data bus is only 16 bits wide. Both of these features are inefficient in the context of the AOSP. We did not consider it a drawback that the 68010 requires a coprocessor for floating-point operations, provided that the coprocessor could be made to operate at sufficiently high speed.

We also examined and rejected the following commercial and aerospace industry architectures on the basis of one or more of the criteria previously described: the National Semiconductor 16032, the Texas Instruments 9900 series, the Rockwell International AAMP (Advanced Architecture Microprocessor), and the Intel iAPX 432 (Table 1). These microprocessors shared a principal drawback—their large gate counts.

Military architectures. We also considered a GaAs implementation of the Mil-Std-1750A[2] architecture because of the large amount of software written for this processor. Because this design is a DoD standard, there was considerable initial enthusiasm within the DARPA community for its selection as the target architecture for the GaAs microprocessor project. However, Mil-Std-1750A specifies a 16-bit architecture with a single 16-bit bus, yielding only 128K bytes of DMA, half of which is used for instructions and half for data. A memory management unit can be appended to the processor in most versions to expand the DMA capability to 2M bytes. However, this approach is not efficient for the large number of array-oriented and I/O operations required in the AOSP. Most implementations do support on-chip floating-point operations and also contain an efficient interrupt handler. Sixteen general-purpose, 16-bit registers are also provided. A very large number of instructions and addressing types are available in this architecture, yielding a variety of instruction formats and lengths. Control logic for the 1750A is implemented with microcode. One implementation of Mil-Std-1750A by Texas Instruments[3] requires two chips, each containing over 16,000 gates and 110K bits of microcode ROM.

We rapidly concluded that Mil-Std-1750A could not provide a suitable GaAs microprocessor architecture for the AOSP project. The processor architecture is too complex and yet does not implement a 32-bit machine. Computation with 64-bit operands is not supported, and the direct address space is considered too small, given the trends toward ever-increasing direct memory capacity.

We also examined the Nebula machine (Mil-Std-1862B), a DoD 32-bit architecture, for implementation in GaAs (Table 1). However, this architecture is not necessarily a single-chip microprocessor, even in silicon VLSI. One of the multichip implementations



attempted to date by Raytheon Corporation requires approximately 44,000 gates, which is much too large for a GaAs version, at least through the end of the decade. The instruction execution rate of this machine is also comparatively slow at 500 KIPS. Therefore, we also eliminated this architecture from further consideration.

RISC architectures. The family of reduced instruction set computers (RISCs) appeared to match the constraints of a GaAs microprocessor more closely than the machines described previously. This type of computer, developed independently by several universities and corporations, was considered a novel approach in the early 1980's. Only recently has it become commercially available in silicon. The family of RISC architectures was developed after it was recognized[4] that a large number of complex instructions normally implemented in hardware on a microprocessor are rarely used in compiled or manually prepared assembly-language code. These infrequently used instructions could therefore be more efficiently implemented in software, reducing the hardware support to a small set of simple instructions.

This design approach can reduce the gate count of the processor considerably, especially in the control section, and can permit the remaining simple instructions to execute much more rapidly. The simplified hardware instruction set complicates the low-level software required for the microprocessor, but this price is paid only once at compile time, rather than every time an instruction is executed. Compile-time penalties are not a major concern if overall performance is improved and the simplified microprocessor exhibits a considerably reduced gate count. Several versions of RISCs had been introduced at the time we performed this architecture study, including the RISC II processor design by the University of California at Berkeley, the 801 minicomputer design by IBM, the transputer design by Inmos Ltd., and the MIPS (Microprocessor without Interlocked Pipeline Stages) design by Stanford University (Table 1). As we will discuss, this last architecture appears to be best suited for implementation in GaAs.

The major strengths of a RISC include a small number of simplified instructions, a reduced gate count, and faster instruction execution time. Several other architectural aspects naturally lend themselves to straightforward hardware implementation on this type of machine (although not every RISC employs all of these features). These features include:

- a load/store architecture (only load and store instructions access memory, while all others operate on registers);
- single-cycle execution of most instructions;
- a short critical path;
- a hardwired control section (rather than microcode);
- a Harvard architecture (separate memories and buses for data and instructions);
- a fixed-instruction format (all instructions are the same size and structure);
- preprocessing of pipeline interlocks in software; and
- the ability to keep major resources of the machine (the ALU, the data memory, and the instruction memory) fully occupied most of the time.

The RISC II. The RISC II[4] contains a large on-chip register file used for storing instructions, local variables, and constants. The register file consists of eight overlapping windows of 32 words each (a powerful support structure for procedure calls). However, the file must be loaded one word at a time. When this machine is interrupted, the contents of the entire register file must be stored in off-chip memory—a very time-consuming operation. The machine supports only 31 instructions in hardware, each of which is one word (32 bits) in length. All instructions execute in one clock cycle, except load or store operations, which require two cycles. The RISC II contains a three-stage pipeline. The compiler targeted for this machine rearranges the instructions to allow useful operations to be executed during pipeline conflicts. The RISC II designers intended that users would almost always program in high-level languages (HLLs), thereby allowing this associated compiler to optimize the code before execution.

Because only a small portion of the chip area is devoted to control, the floor plan of the device is very regular. The addresses and data words are multiplexed to reduce the total I/O count for the chip to 50 pins, but the multiplexing degrades I/O operation speed. Programs written for the RISC II are approximately 50 percent larger, but typically run faster, than programs written for the more conventional machines we described earlier. Floating-point operations are handled off chip, and integer multiply and divide operations are executed in software. Interrupts are supported only with an external interrupt flag; all other interrupt operations must be processed off chip. The inability of the RISC II to process interrupts on chip is a significant limitation in the AOSP application environment.

The complexity of the RISC II is 12,000 gates, which could have been reduced somewhat in a GaAs version by decreasing the size of the register file at the cost of a severe performance degradation. Although the RISC II offered many attractive features, it did not appear to be efficient for the type of application typical of the AOSP environment, and therefore we did not recommend it for implementation in GaAs.

The IBM 801 machine. The 801 minicomputer designed by IBM[5] is a RISC in which the most recently used instructions are stored in an instruction cache which has an access time of one machine cycle. A data cache is also provided. Data from the two caches are fetched asynchronously to one another, and it is possible to overlap access to the two caches. A software I/O manager synchronizes the caches when required, minimizing the execution of unnecessary load and store operations. A compiler reorders the code to allow useful instructions to be executed during branch and load delays. This processor was also intended to be an HLL-computer, using a very efficient compiler to produce object code nearly as efficiently as the best hand-generated code.

Thirty-two general-purpose, globally allocated registers are available in the 801 machine, so that data needed again within a short duration is available immediately. Instructions are 32 bits in length and have been optimized for microcode. Complex instructions are executed in software. Addresses and data values are also 32 bits in length. A fixed-precision multiplication can be completed in 16 cycles, while a division can be completed in 32 cycles. Floating-point operations are executed off chip with the 801. This machine averages 1.1 cycles per instruction, exhibits cache hit ratios close to 100 percent, and contains a two-stage pipeline. The number of instructions required for a program varies considerably. In some cases, instruction streams are equivalent in length to those found in a more complex processor. In other cases, a 50 percent increase in the number of lines of code is required, particularly for applications requiring many floating-point operations. The information to determine if this machine has an acceptable gate count for a GaAs implementation was not available. The architecture is proprietary to IBM, and additional information is not accessible by the public. This machine appears to match the AOSP environment and might be applicable to a variety of tasks in GaAs if the design were made available.

The Inmos transputer. The transputer is a very high speed 32-bit microprocessor executing a reduced-instruction set of 59 instructions.[6] The current implementation requires approximately 62,500 gates, including 4K bytes of SRAM, a memory interface, a peripheral interface, and on-chip serial links to other transputers, as well as the main processor. A transputer can be used alone, although the intention as stated by Inmos is to employ networks of transputers to increase the performance of an entire system. The same software can be used regardless of the number of processors used. The transputer, which responds very quickly to external interrupts and supports simultaneous block transfers, was designed to implement concurrent processes. Instruction sequences are executed with no wait states.

The MIPS machine was the only architecture which satisfied all constraints of a GaAs microprocessor for the AOSP.

The transputer can be programmed in most high-level languages, but is intended to perform most efficiently when programmed in Occam, a language specifically developed by Inmos to exploit concurrency. Occam eliminates the need for assembly-language programming and is the lowest level language supported by the transputer. Floating-point operations are executed in software in the early versions of the transputer, and, we believe, are to be executed in hardware in the later versions. The transputer supports a single 32-bit bus, which must be multiplexed for data and addresses. The transputer employs an internal microengine to execute its assembly-level instructions, and possesses no completely general-purpose registers in the hardware. These features are unusual in a RISC. The transputer is also a proprietary architecture. If this design had been available, we might have considered it as a replacement for the entire NCU portion of an AOSP ACE (implemented with multiple chips), rather than merely as the control processor in the NCU.

The MIPS machine. Stanford University originally developed the MIPS machine[7-10] in 1981 as a RISC-architecture project under DARPA sponsorship. After examination, the MIPS machine was the only architecture which satisfied all constraints of a GaAs microprocessor for AOSP. The philosophy behind the MIPS project was to implement in hardware only the most frequently executed and time-consuming instructions and to implement infrequently used instructions in software. This implementation of the instructions increased the overall speed of operation of the processor. The MIPS machine architecture and supporting software were available for government-supported projects. The MIPS is a very high speed 32-bit architecture, primarily due to an efficient match between its architecture and its supporting software.

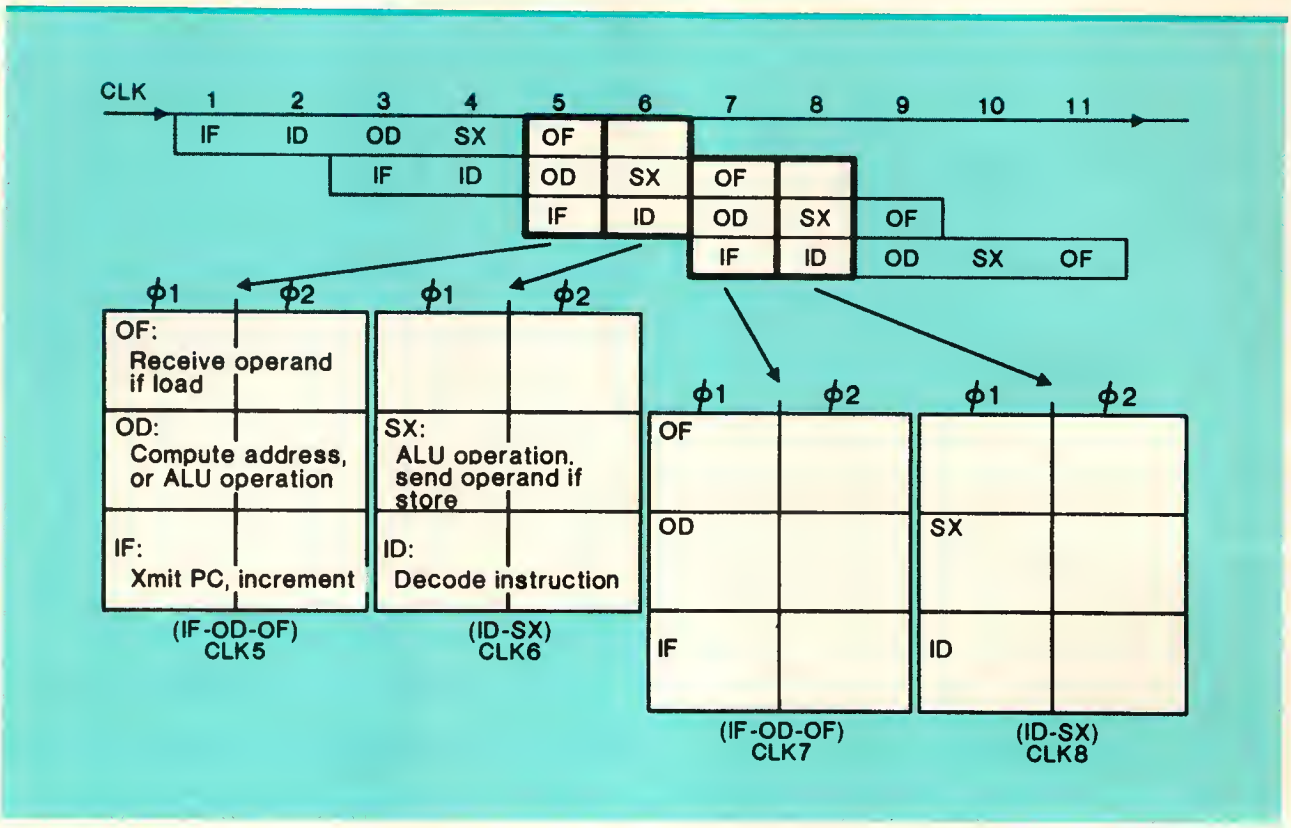


Figure 3. Stanford MIPS five-stage pipeline with active pipe states for each clock cycle. Reproduced in part from Hennessy et al.[9]

The Stanford MIPS machine contains a five-stage pipeline, with a new instruction entering the pipe every other cycle, yielding three active instructions per cycle. Each pipestage executes in the same amount of time, and every machine instruction passes through each stage. The five stages are called

- *IF* (instruction fetch), which transmits and increments the program counter for fetching the instruction;
- *ID* (instruction decode), which decodes the instruction;
- *OD* (operand decode), which either computes the memory address for a load or store instruction, computes the program counter for a branch instruction, or performs an arithmetic operation;
- *SX* (store and execute), which transmits the operand for a store instruction, and, in addition, either performs an arithmetic operation or performs the compare for a conditional instruction; and
- *OF* (operand fetch), which receives the operand for a load instruction.

The machine alternates between the IF-OD-OF cycle and the ID-SX cycle, as long as no interrupts occur, as shown in Figure 3. This pipeline implementation enables 100 percent utilization of the major hardware resources of the processor, which include the instruction memory (by the IF and ID pipestages), the ALU (by the OD and SX pipestages), and the data memory (by the OF and SX pipestages), as depicted in Figure 4. Because two pipestages are able to execute arithmetic operations, two adds can be packed into one machine instruction and executed within one machine cycle. A 32-bit fixed-precision multiply can thus be executed in eight cycles, with a 2-bit Booth's algorithm executed at 4 bits per cycle. An ALU operation can also be packed with a load into one machine instruction in this architecture.

Interlocks are necessary in pipelined microprocessor architectures to prevent instructions from interfering with one another as they traverse the pipeline. The MIPS machine implements its pipeline interlocks in software, rather than in hardware as is customary in other microprocessor designs. This implementation eliminates a significant amount of complex control logic from the processor. Interlock arbitration can be accomplished in software because the information necessary to generate the interlocks is known at assembly time. By moving this function from execution time to compile time, execution time becomes much faster. The MIPS uses a sophisticated reorganizing assembler which is tightly coupled with the simpler hardware. This "reorganizer" assembles the code into executable machine code, packs two independent instructions into one instruction where possible, and reorganizes the code. Code reorganization accomplishes pipeline interlocking by substituting useful instructions from elsewhere in the code stream where No-op instructions would otherwise have to be inserted to avoid pipeline conflicts, as shown in Table 2. Reordering of instructions improves code execution speed an average of 20 per-

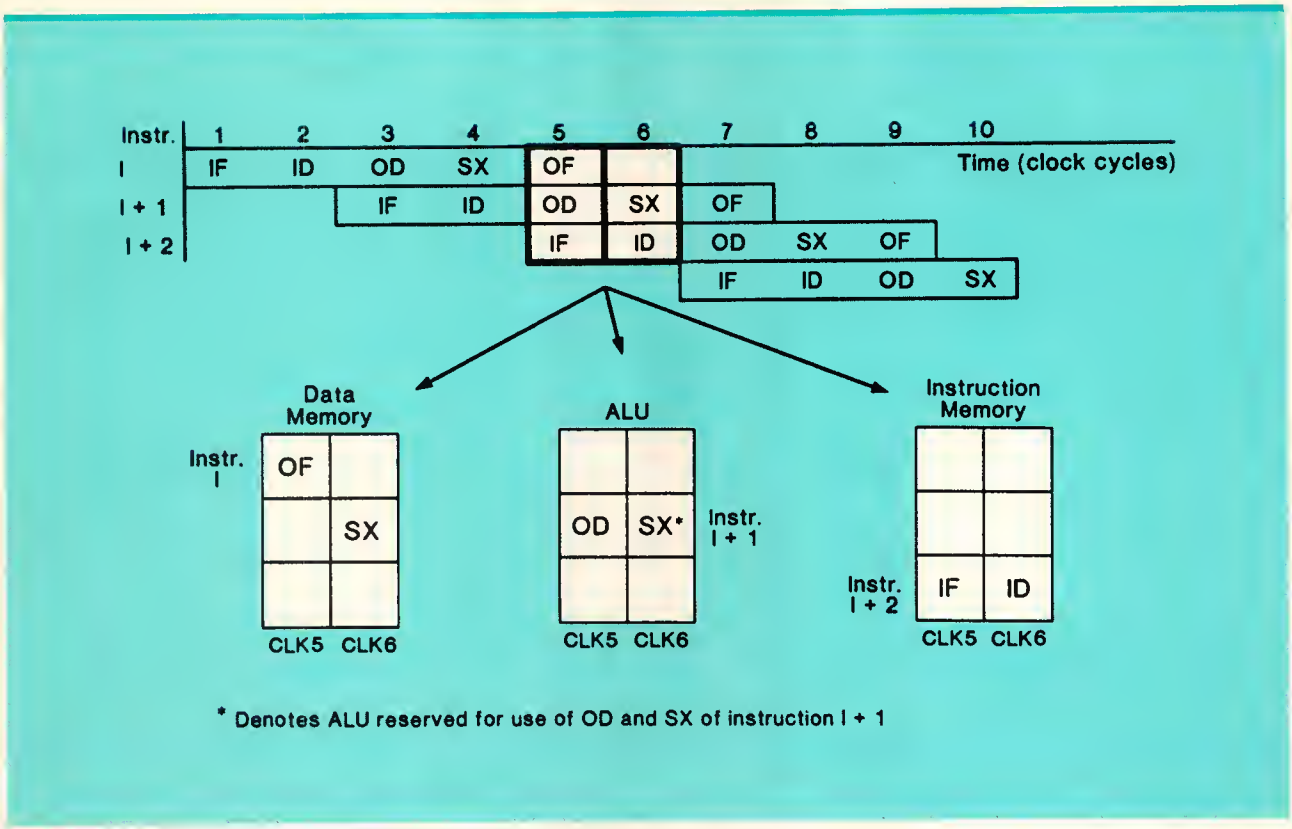


Figure 4. Stanford MIPS five-stage pipeline showing hardware resource utilization by each pipestage. Reproduced from Przbylski et al.[7]

Table 2.
Example of MIPS code before and after reorganization and packing by the sophisticated assembler.*

	Source code	Correct code with No-ops	Reorganized code
(For Loop) →	Begin		
	A[i]:=B[i]+C[i];	L20: ld (r4, r1), r6	L100: ld (r4, r1), r6
		ld (r5, r1), r7	ld (r5, r1), r7; add r6, r8
		No-op	
		add r7, r6, r9	add r7, r6, r9; add r7, r10
		st r9, (r3, r1)	st r9, (r3, r1); add #1, r1
	R:=R+B[i];	add r6, r8	
	S=S+C[i];	add r7, r10	
		add #1, r1	
		ble r1, r2, L20	ble r1, r2, L100
		No-op	st r8, R(sp)
		No-op	st r10, S(sp)
	End	st r8, R(sp)	
		st r10, S(sp)	

*Note that the code length has been reduced by almost half. Reproduced from Hennessy et al.[10]

cent.[7] Instruction packing increases code density, and reduces the execution time of an instruction stream by an additional 30 percent. This gives a combined improvement in execution speed of over 50 percent, when compared to the execution speed of code produced by traditional assemblers. Because the reorganizer accepts MIPS assembly code as input and produces machine-executable code as output, this software module must be unique to each hardware implementation.

The Stanford MIPS machine implements conditional control flow with software by using a compare and branch instruction, which executes in a single machine cycle. The MIPS does not use condition codes—the typical method of implementing conditional control flow—because the structure of the hardware typically employed for condition codes is irregular, both in design and in physical layout. Thus, condition codes are an inefficient use of space on the chip. A total of 16 comparisons, both signed and unsigned, are available in the ALU for this purpose. The MIPS also implements a "set conditional" instruction with the same 16 comparisons.

The Stanford MIPS machine employs a Harvard architecture, having separate data and instruction memories which are accessed on alternating phases of the two-phase clock. Only load and store instructions access memory; all other instructions operate on registers (thus referred to as a load/store architecture). Five addressing modes are available: long-immediate, absolute, based, indexed, and base-shifted by n (where $0 < n < 5$). Although the Stanford MIPS is a word-addressed machine for both instructions and data, it has special instructions to support byte operations on data. Data is accessed with a 24-bit physical address, expandable to a 32-bit virtual address. This yields 16M words of directly addressable memory, with one level of optional page mapping off chip. In addition, if an instruction is *not* a load or store (approximately 50 percent of the instructions), a memory cycle may be used for cache write-back or prefetch.[7]

The Stanford architecture supports vectored interrupts with a 12-bit address concatenated to the current machine status for a jump to the proper interrupt routine. When any type of exception is detected, instructions that are currently in the pipeline are completed, if possible. The state of the processor is stored in the "surprise register." Execution is transferred to memory location zero, where the registers and three return addresses are saved for return to the code sequence that was in progress, and the interrupt handling routine is invoked.

The original MIPS chip was fabricated in NMOS and requires 8000 gates and 84 I/O pins. The chip layout is very regular, containing six major hardware sections. The control of the microprocessor is divided into two functional units implemented as PLAs: (a) the instruction decode unit, and (b) the master pipeline control. These occupy about 20 percent of the chip area. The data path is interconnected through a pair of 32-bit buses. It consists of the ALU, a barrel shifter, a register file (including 16 general-purpose registers), and the program counter and address mask unit. A control bus interfaces the data path to the control sections.

These chips have an average throughput rate of two million instructions per second and a power dissipation of 1.6W, at a clock rate of 4 MHz.[9] The Stanford MIPS was benchmarked against a Motorola 68000 and a VAX 11/780 by executing several computation-intensive programs written in Pascal. The results of these benchmarks demonstrated that the average throughput of the MIPS was five times faster than that of the Motorola 68000 and twice as fast as that of the VAX 11/780[8] (Table 3).

The MIPS instruction set originally defined at Stanford[7] contains 32 instructions. All of these instructions are 32 bits in length, have the same instruction format, and execute in a single machine cycle. The Stanford MIPS machine requires a coprocessor for high-speed execution of floating-point operations. However, if one is not attached, these instructions can be executed in software on the main processor.

Current projects in GaAs based on the MIPS architecture

After completing our 1983-84 studies, we felt that an entire MIPS-type microprocessor might be implemented as a custom-designed GaAs IC containing approximately 10,000 equivalent gates, because of the relatively low complexity of such a microprocessor. We also felt that a custom-designed floating-point coprocessor might be implemented in an additional 10,000 gates on a second chip.

Based on these results, in 1984 DARPA initiated a five-year program to develop a full 32-bit architecture microprocessor and floating-point coprocessor in gallium arsenide using the Stanford MIPS machine as the baseline. Three contractors—McDonnell Douglas, Texas Instruments teamed with Control Data Corporation, and RCA Corporation teamed with TriQuint Semiconductor—completed a one-year architecture study phase. A four-year, two-phase project to fabricate a 32-bit GaAs microprocessor is currently under way, both at McDonnell Douglas and at Texas Instruments. The goals for these GaAs microprocessors are

- a custom main processor chip and a custom floating-point coprocessor chip,
- a chip complexity of 10,000 gates,
- operation of both chips at a 200-MHz clock rate, and

Table 3.
Results of Pascal benchmarks in seconds.*

	MIPS	M68000	VAX 780 (appr.)	DEC 20/60
Clock speed	(4 MHz)	(8 MHz)	(5 MHz)	
Si transistor count	(25,000)	(65,000)		
Puzzle	2.40	6.1	5.2	2.6
Queen	0.44	1.9	1.0	0.5
Perm	0.56	3.0	1.2	0.6
Towers	0.64	2.9	1.4	0.7
Intmm	0.80	5.0	1.0	0.5
Bubble	0.58	3.7	1.4	0.7
Quick	0.41	2.6	0.8	0.4
Tree	1.01	9.9	2.0	1.0
Avg. (relative to MIPS)	1.00	5.1	2.0	1.0

*Reproduced in part from Przbylski et al.[7]

- a sustained throughput rate of at least 100 million instructions per second.

The first phase ended in March of 1987 with each of the companies producing

- the assembler, linker, and reorganizer software;
- stand-alone demonstration chips containing large portions of the microprocessor; and
- a more detailed system-level specification.

During the second phase, the actual chips and boards will be fabricated by the two contractors, and the systems will be integrated into single-board computers. DARPA plans to receive the first prototype GaAs RISC CPU chips and floating-point coprocessor chips by early 1988. DARPA then expects completion of the first GaAs single-board computers for the GaAs AOSP project in early 1989.

A computer operating at a 200-MHz clock rate presents a serious design problem: data starvation. The main memory currently cannot supply instructions and data operands to the processor at a sufficiently high rate to keep it occupied. The provision of fast cache memory helps alleviate this problem. Both GaAs RISC development projects plan to use two off-chip cache memories (one for instructions and one for operands) and two cache controllers or memory manager chips.[11,12] The caches will be necessary in both the autonomous single-board computer system and the GaAs AOSP. These caches will be 1K words each, with an optimum access time of one nanosecond and a worst case access time of 2.5 ns. The complete computer system will contain

- the RISC CPU,
- two floating-point coprocessors,
- two cache memories,
- two memory management units,
- a system controller for low-speed I/O and external exceptions, and
- main memory.

The CPU chip may also function as an embedded controller for a system which supplies its own memory and does not require the floating-point coprocessor.

We believe that there is also considerable growth potential in the GaAs version of a MIPS processor chip set, which could be explored as follow-ons to the present projects. For example, with more available real estate on the chips, a cache memory or additional registers could be placed on chip. The addition of on-chip cache would provide a significant performance gain, since in current GaAs RISC designs a transfer bottleneck occurs between the microprocessor and its off-chip caches. Alternately, some of the floating-point functions now performed on the coprocessor chip could be incorporated into the microprocessor itself.

In addition, Rockwell International preliminary studies indicate the possibility of a MIPS-based machine operating at a 500-MHz clock rate and executing more than 200 million instructions per second. This version of the MIPS machine would require second-generation GaAs transistors, such as HEMTs or HBTs.

Silicon versions of the DoD MIPS machine. The GaAs RISC project has had an impact on concepts within DARPA, the US Air Force, and the Strategic Defense Initiative Office regarding optimum

architectures for next-generation processors. In order to gain additional benefits from GaAs RISC developments in architecture, hardware, and software (see discussion on software), DARPA has initiated a parallel development in silicon. These microprocessors will be fabricated using bulk CMOS and CMOS/SOS to exploit the higher complexity levels available in silicon. Three contractors—Sperry, General Electric, and RCA—completed a one-year architecture study phase in 1986. Two of these contractors are now conducting a two-year phase to fabricate a silicon MIPS-type chip containing on-chip cache, and a floating-point coprocessor chip, with the target clock rate of 40 MHz. Silicon chip prototypes were successfully demonstrated by the two contractors in the late fall of 1987.

Establishment of a standard ISA and transportable software. The government agencies involved with these projects decided to establish a standard Instruction Set Architecture (ISA) for all MIPS-based machines. This standard provides software compatibility for all MIPS-based processors presently under development while still allowing flexibility at the hardware design level. An ISA is a list of attributes visible to an assembly-language programmer or to an HLL compiler. In general, an ISA description includes lists of data formats, instruction formats and mnemonics, addressing modes, available registers, memory and interrupt control structures, I/O operations, and detailed instruction-set requirements. However, an ISA does not include specific details for a given hardware implementation of a processor.

Carnegie Mellon University, Stanford University, Mayo Foundation, and the contractors currently implementing MIPS-based hardware collaborated to develop an ISA document for the DoD MIPS machines.[13] Their goal was to ensure that all versions of the MIPS processor will execute the same software on an interchangeable basis. This MIPS-based ISA currently serves as the baseline document for HLL compilers (in Pascal and Ada) under development through DoD sponsorship for the MIPS processors. Translators, or cross-assemblers, are also being written for the Motorola 68000 and for the Mil-Std-1750A microprocessors. Assembly-language programs previously written for these processors can then be used directly on any of the DoD-sponsored MIPS processors. This software is now available, with the exception of the Ada compiler. The Software Engineering Institute associated with Carnegie Mellon University is maintaining this software.

We now briefly describe several features required by the DoD-standard ISA for all current and future MIPS-based processors. The ISA increased the number of assembly-level instructions from 32 to 69. The additional instructions support

- unsigned arithmetic operations;
- byte, half-word, and word operations; and
- floating-point operations.

Most of these added instructions should still execute within a single machine cycle. Floating-point operations are executed in hardware on the coprocessor, or—in the absence of a coprocessor in applications not requiring high-speed floating-point operations—in software on the main processor. Single-precision integer and floating-point data operands are 32 bits wide, while double-precision data operands are 64 bits wide. At least 16 general-purpose, 32-bit registers and at least four 64-bit, floating-point registers must be available on these processors to satisfy the ISA. All data memory access is by explicit load and store instructions. Instructions continue to be addressed on word boundaries. However, addressing of data operands has been changed to a byte-addressed format. As a result of this change, data can be loaded and stored as byte, half-word, and word-sized operands. Only two addressing modes have been specified in this ISA: *absolute*, in which the address is specified directly; and *based*, in which the address is obtained by adding an offset to a base register.

Several changes not specified in the standard ISA have also been made, which appear to improve the efficiency of the GaAs implementations of MIPS in comparison to the original Stanford MIPS architecture. These changes include

- modification of the number of pipestages from five to three in one implementation and to six in the other;
- initiation of a new instruction every machine cycle, rather than every other cycle; and
- elimination of the capability of packing two assembly-language instructions into one machine instruction.

(Instruction packing unnecessarily complicates the design of the processor control section.)

An additional feature of the GaAs RISC implementations is that up to eight coprocessors can be attached to a single microprocessor chip. These coprocessors do not necessarily have to be floating-point coprocessors.

Coprocessor implementations. Because the coprocessor architectures currently being implemented with the GaAs RISC processors vary considerably, we present only one of the coprocessor implementations under development.[11] The microprocessor and its coprocessor operate synchronously from a common clock and have common instruction address, instruction data, and operand data buses. The microprocessor calculates memory addresses and initiates memory references and then continues executing its own instructions. The coprocessor accepts its own instructions and data and outputs its values either to mem-

GaAs RISC implementations allow eight coprocessors to be attached to a single microprocessor chip.

ory or to the processor. The system presently is designed to operate with two floating-point coprocessors attached to a single processor. All three operate concurrently. The design also allows for the addition of up to six additional coprocessors. Each of the two coprocessors has a separate operation code field in the instruction word. The coprocessor interrupts the processor on a floating-point exception. The microprocessor then determines the cause of the exception by reading the coprocessor status register, and the appropriate software handler is executed by the microprocessor. The coprocessor monitors the processor's status bits and uses this information to track the pipeline-stage sequencing in the processor to detect wait states and exceptions that may affect the operations within the coprocessor.

All coprocessor instructions have a fixed execution time with no operand dependencies. This simplifies the reorganizer's scheduling task. The predictability of coprocessor operations and the presence of instruction prefetching simplify pipelining. Instruction prefetching means that a new instruction word is fetched on the last cycle of the previous instruction. Two arithmetic instructions are loaded into the coprocessor with one instruction word, so that the second instruction can be initiated immediately upon completion of the first. The control signals for the operating cycles of each coprocessor are decoded one cycle ahead, eliminating control decode time and making the cycle time dependent only on the data path.

This coprocessor design is divided into two sections—the bus interface unit and the arithmetic unit—each with separate control. The chip area devoted to control of the coprocessor is very small, consisting of two small PLAs and using only 20 percent of the transistors on the coprocessor. The bus interface unit performs conditional branching and testing and monitors the instruction bus while the arithmetic unit is processing data. This architectural approach allows the CPU to execute a "branch on coprocessor busy" operation rather than executing No-op instructions. This approach also enables a branch based upon the results of the first arithmetic instruction while the second instruction is executing in the coprocessor.

The arithmetic unit is optimized for floating-point operations and provides full double-precision data paths. Eight double-precision operand registers are also available. Operands are converted to single-precision values only during load and store operations. The exponent and significant processors of the arithmetic unit are separate and operate in parallel. A single- or double-precision add or subtract operation requires four clock cycles to complete, plus two additional clock cycles if normalization is required. The multiply operation is performed 2 bits at a time, requiring 15 cycles to complete a single-precision multiplication and 30 cycles to execute a double-precision multiplication. The coprocessor hardware supports four rounding modes from the *ANSI/IEEE Standard 754-1985 for Binary Floating-Point Arithmetic*, as well as six exceptions.

This coprocessor implementation appears to be very efficient and should demonstrate a substantial execution rate.

Silicon microprocessor designers have capitalized on the availability of large numbers of gates on VLSI chips to create complicated architectures based on parallelism and a complex set of assembly-language instructions. A gallium arsenide implementation of a microprocessor requires an alternate design approach with the availability of a limited number of gates and very fast transistors. Our study of microprocessor architectures demonstrated that a simple design, implementing in hardware only a small set of frequently used assembly-language instructions, is better suited to a GaAs fabrication technology than is a more complex design. More complicated functions are implemented in software tightly coupled to the hardware. Several computer architectures with these characteristics, known as RISCs, have been developed in the past few years. A review of these architectures indicated that one of them, the MIPS machine developed by Stanford University, could be fabricated in GaAs with approximately 10,000 gates and may operate at a clock rate of 200 MHz with a sustained throughput of 100 million instructions per second. Digital GaAs technology can take advantage of the fact that the MIPS architecture is very simple and relies heavily upon an instruction pipeline and a sophisticated assembler.

The original goal of our microprocessor evaluation project was the identification of a 32-bit microprocessor architecture that could not only execute at high speed but could also be entirely implemented on a single GaAs chip (with the exclusion of the floating-point coprocessor) in the near future. Because the MIPS architecture appeared to offer this possibility, two GaAs implementations of this architecture are presently under development,[11,12] as is the required software support. The DoD-sponsored MIPS will be able to function both as a stand-alone computer and also as a microprocessor embedded within more complex processors fabricated either in silicon or GaAs, as in the DARPA/Air Force AOSP.

Acknowledgments

This research was supported in part by contracts MDA-903-84-C-0324, N66001-85-C-0337, and F29601-84-C0016 from the Defense Advanced Research Projects Agency. The authors wish to thank C.L. Bates, D.D. Endry, S.M. Karwoski, L.M. Krueger, J.M. Ryan, M.L. Samson, D.J. Schwab, B.R. Shamblin, R.L. Thompson, C.R. Treder, T.L. Volkman, and S.K. Zahn of the Mayo Foundation for technical assistance; S. Roosild of DARPA and S. Karp for helpful discussions; and E.M. Doherty and S.J. Richardson of the Mayo Foundation for preparation of text and figures.

References

1. B.K. Gilbert et al., "Signal Processors Based Upon GaAs ICs: The Need for a Wholistic Design Approach," *Computer*, Oct. 1986, pp. 29-43.
2. *Military Standard Sixteen-Bit Computer Instruction Set Architecture*, MIL-STD-1750A (USAF), US Gov't Printing Office: 1980-603-121/3202, Department of Defense, Washington, DC, May 21, 1982.
3. *VHSIC Data Processor Unit Integrated Circuit Specification*, Preliminary Draft, Revision C, Texas Instruments, Dallas, Tex., Feb. 21, 1985.
4. M.G.H. Katevinis, *Reduced Instruction Set Computer Architectures for VLSI,* doctoral dissertation, Univ. of California, Berkeley, Calif., Report No. UCB/CSD 83/141, Oct. 1983.
5. G. Radin, "The 801 Minicomputer," *IBM Journal of Research and Development,* May 1983, p. 237.
6. I. Barron et al., "Transputer Does 5 or More Mips Even When Not Used in Parallel," *Electronics,* Nov. 17, 1983, p. 109.
7. S.A. Przbylski et al., "Organization and VLSI Implementation of MIPS," Tech. Report No. 84-259, Computer Systems Laboratory, Stanford Univ., Stanford, Calif., April 1984.
8. J.L. Hennessy et al., 'Hardware/Software Tradeoffs for Increased Performance," Tech. Report No. 228, Computer Systems Laboratory, Stanford Univ., Feb. 1983.
9. J.L. Hennessy et al., "Design of a High Performance VLSI Processor," Tech. Report No. 236, Computer Systems Laboratory, Stanford Univ., Feb. 1983.
10. J.L. Hennessy et al., "MIPS: A VLSI Processor Architecture," Tech. Report No. 223, Computer Systems Laboratory, Stanford Univ., June 1983.
11. T.L. Rasset et al.,"A 32-Bit RISC Implemented in Enhancement-Mode JFET GaAs," *Computer*, Oct. 1986, pp. 60-68.
12. E.R. Fox et al., "Reduced Instruction Set Architecture for a GaAs Microprocessor System," *Computer*, Oct. 1986, pp. 71-80.
13. T. Gross and R. Firth, *Core Set of Assembly Language Instructions for MIPS-based Microprocessors*, Version 3.2, Software Eng. Inst., Pittsburgh, Pa., Nov. 18, 1986.



Barbara A. Naused has been a computer design engineer at the Mayo Foundation's special-purpose processor development group since 1981. Her experience includes signal processor architecture systems and GaAs integrated circuit design, comparative reviews of microprocessor architectures, and hardware and software design.

Barbara Naused received her BS degree in applied mathematics from the University of Wisconsin in 1981 and is currently completing an MS degree in electrical engineering from the University of Minnesota.



Barry K. Gilbert is staff scientist and director of the special purpose processor development group at the Mayo Foundation. He is currently responsible for the design and development of computer-aided engineering tools and hardware design methods to exploit the speed performance of GaAs digital ICs in the next generation of high-performance signal processors. His interests include the development of algorithms for the real-time analysis of wide-bandwidth image and signal data and the design of specialized computers to execute these tasks.

Barry Gilbert received his BS degree in electrical engineering from Purdue University in 1965 and his PhD in physiology/biophysics from the University of Minnesota, Minneapolis, in 1972.

Questions about this article may be directed to Barbara Naused at the Department of Physiology and Biophysics, Mayo Foundation, Rochester, MN 55905.

Reader Interest Survey

Indicate your interest in this article by circling the appropriate number on the Reader Interest Card.

Low 150 **Medium** 151 **High** 152

FEATURE

A Distributed System for Real-Time Applications

Eli T. Fathi, Applied Silicon Inc. Canada
Eloi Bosse and Jean Caseault,
Communications Research Centre, Canada

This versatile network uses multiple microprocessors and a split-bus architecture to collect and process real-time data from a variety of sources and then transfer it to different destinations.

Modern communication applications need computer system architectures that can collect, process, and distribute large amounts of digital data. The amount of data in the system directly impacts the type of processing the system can handle as well as the system bandwidth. The handling of the digitized data (namely the collection, preprocessing, combining, postprocessing, and distribution of this data) requires complex automated procedures and special-purpose architectures.

To acquire and process data in an effective fashion, one must develop a versatile distribution network that is transparent to the system computer architecture. Such a network must not interfere with the natural flow of data in the system but rather support it throughout the various stages, from initial acquisition through postprocessing operations.

System functional requirements

The underlying motivation behind the development of the distribution system architecture we describe was our need for a general-purpose testbed for radar applications. The most important factors influencing the design of this system are the constraints imposed by the radar application itself and the current state of microprocessor technology. The radar system produces large amounts of information coming from different sources (refer to Figure 1). All this information must be directed simultaneously to multiple destinations.

The volume of data, and the type of data processing, made it impossible for any currently available (at the time the design was undertaken), single-chip microprocessor to handle the complete preprocessing load alone. We considered the possibility of using bit-sliced LSI devices, with a microprogrammed control store and lookup tables, to provide the first level of data processing. However, we decided on a more comprehensive, general-purpose solution that splits the overall processing requirement into a number of well-defined functions (for example, doppler filtering, averaging). Dedicated hardwired logic could then be used to execute these functions. In addition, it was possible to use a powerful processor to aid in the computation of lengthy functions. Additional processing power might be needed mainly to coordinate the various activities.

The system must be capable of processing large blocks of radar data in real time. It must therefore be powerful, fast, and versatile to meet the experimental requirements. It must also be capable of being easily modified or enhanced. In considering its implementation, we investigated both single processor- and multiprocessor-based system designs. We found the

0272-1732/87/1000-0021$01.00 © 1987 IEEE

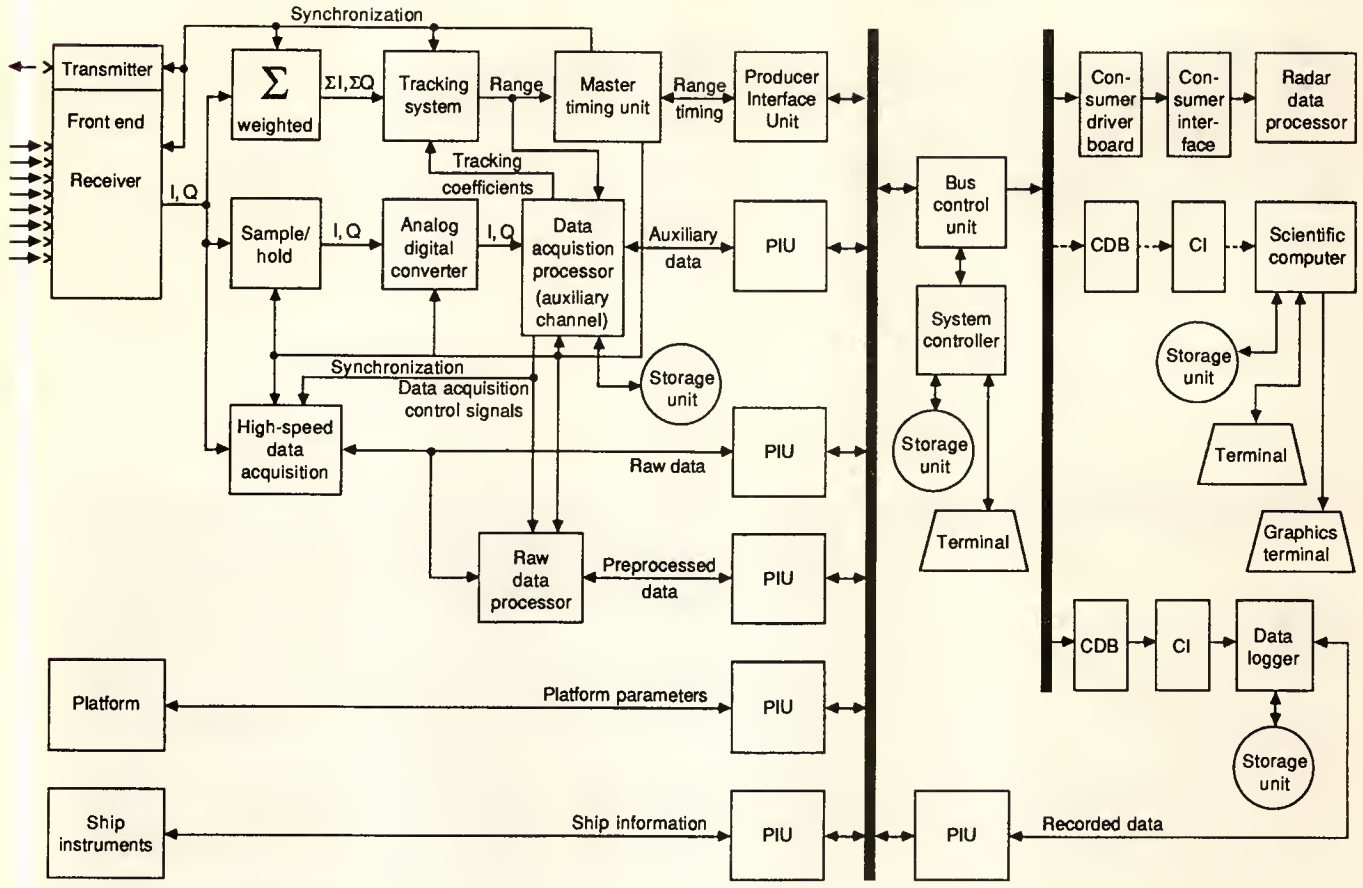


Figure 1. System architecture.

uniprocessor approach unattractive because of its poor cost-to-performance ratio, its inflexibility, its inability to accommodate future changes and/or enhancements, and the low reliability factor associated with a single device.[1] We identified the multiprocessor approach as the most suitable choice for meeting the most critical requirements of the application.

To provide the real-time processing capability and to maintain system flexibility, we chose a multilevel tree structure as the basic computer architecture. This architecture was decomposed into three major sections: the preprocessor section, the distribution network, and the postprocessor section. By decomposing the system in this fashion and providing the appropriate infrastructure, one can alter the functionality of the system by changing either the front-end or back-end portions. For example, the radar sensor could be replaced by one that operates in a different frequency band. Also, a new special postprocessor could be added in parallel with the existing processor without the need to modify hardware or software.

Some of the inherent benefits associated with this tree architecture are:

- A reduced number of processing elements across the various levels in the tree.
- Decreased software complexity, as units are self-contained, dedicated processing elements, each with its own modular software.
- A potential increase in allocated processing time, due to possible data reduction as data moves through the various levels.
- A reduced number of soft errors accumulated since processing is done near the point of acquisition.
- Fail-soft capability, as real-time diagnostics routines can be performed between adjacent levels. The system is capable of masking out bad data as a result of software errors and/or hardware malfunction.
- Accommodation of new devices such as faster microprocessors and special-purpose LSI/VLSI hardware modules at a later date.

Distribution network architecture

We based the design of the system distribution network upon Anderson and Jensen's study of interconnection transfer strategy, control method, and path structure.[2] In particular, we placed emphasis on studying the system architecture by considering the type of communication allowed between the various modules rather than the structure of the modules themselves.[3]

We developed an interconnection structure that approached the design of the transfer strategy by combining the best features of the direct method (such as is found in single-bus architectures) with the best features of the indirect method (such as is found in star interconnection structures). This hybrid interconnection structure employs a single bus for the input devices and another single bus for the output devices. The use of two separate buses is dictated by reliability considerations,[4] as well as speed of operation. Between the two buses, a bus controller effectively provides a star configuration[5] by supporting centralized routing (transfer control method) with dedicated paths (transfer path structure) to the two buses. We term the system architecture, which is the result of this unique combination of different transfer strategies, a double star configuration.[6]

The whole distribution network can be considered as a loosely coupled multiprocessor system with each processor having its own local program and data memory as well as I/O. In addition, some of the processors share a common communication memory to facilitate the fast transfer of data. By using multiple processors, we reduce the system software complexity significantly. Each processor has its own simple executive to manage the resources and transactions within its domain. Thus, by increasing the amount of hardware, we eliminate the need for a complex real-time, multitasking executive, which would have been required on a uniprocessor system.

The hub of the system is a dual-processor subsystem consisting of a master minicomputer and a slave microprocessor. This dual-processor configuration initiates, monitors, and controls all the transactions occurring within the domain of the distribution network. We refer to the minicomputer master as the system controller, or SC; the specially designed microprocessor slave system is the bus control unit, or BCU. Aside from the SC, the distribution network supports two other types of devices: *producers*, which output data to the distribution network and *consumers*, which remove data from the distribution network.

This interconnection scheme contains two main buses that are physically isolated from each other via a special bus control unit (see Figure 1 again). This physical separation matches well the logical separation between producers and consumers. By following the dual-bus approach, the consumer and producer parts of the system can be expanded independently and selectively. Thus, if additional producing devices must be added to the system, the physical characteristics of the consumer devices bus are not affected and vice versa. Furthermore, the addition of more consuming devices on the bus would require minimal software/hardware changes to the BCU.

Each of the producing devices connects to a producer interface unit, or PIU, to match its own characteristics (serial, parallel, analog) to those of the producer bus (see Figure 2). This permits simultane-

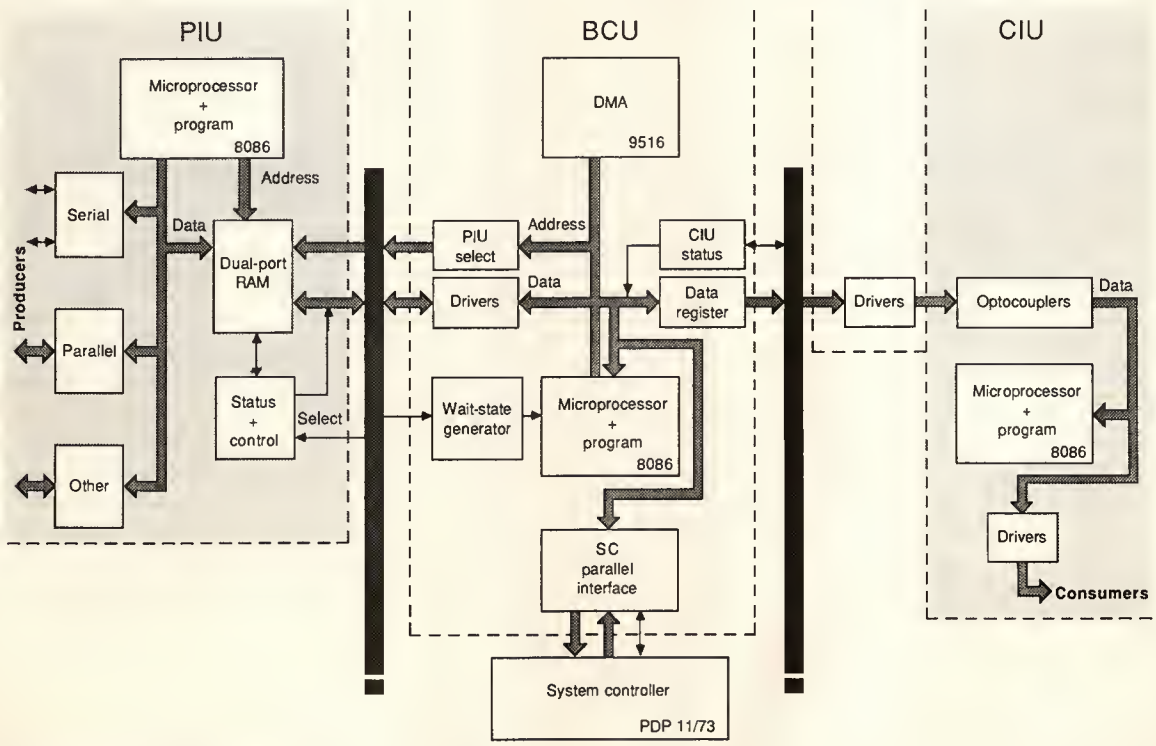


Figure 2. Distribution network.

ous reception of data from each producing device. The transfer of data from the producers to the consumers, and possibly between producers, is initiated and controlled via the BCU. To facilitate this transfer, we designed a bidirectional producer bus.

The consumer bus is a unidirectional, high-drive bus with multiple receivers that are electrically isolated from the producing devices. Any interested consumer has the capability of picking up any or all of the data associated with a given block via its consumer interface unit, or CIU.

System controller

In the proposed configuration, the system controller is the hub of the network; its main control functions are monitoring the system, initiating control functions, and interfacing with the operator. The BCU, acting as a slave to the SC master, controls and monitors the distribution network. The SC communicates with the BCU via a dedicated parallel interface using a single encoded command word. With this command word, the system controller can send or receive information or commands to/from any producer, consumer, or the BCU. The BCU accepts a command, decodes it, executes it, and sends the obtained information back to the SC.

We maintained a very simple communication protocol between the SC and the BCU. While the BCU executes a command, it can be interrupted by, at most, one more urgent command at a time. In addition to the basic command word, a special command word is used by the BCU to notify the SC of an urgent request.

Since all commands must originate either from the SC or be approved by it, the SC, at any one time, has the exact global picture of the whole system. The system operator, through the use of the system monitor, can easily determine the status of the system and control its operation by initiating the appropriate command on the keyboard.

Bus control unit

The SC, being an off-the-shelf minicomputer, lacks the appropriate hardware modularity and software flexibility to manage the complete distribution network by itself. Aside from the requirement for multiple interfaces to multiple sources and destinations, the system sets a stringent requirement on the data transfer rates (a high data transfer rate per device, coupled with data gathering from multiple devices). These conditions made the performance unattainable for a single off-the-shelf device. Therefore, we had to develop a tailor-made subsystem that could accommodate the various devices involved in the transaction, by interfacing to the SC on one side and to the remaining devices on the other. This led us to develop the BCU. We designed the BCU as a slave to the SC, having it execute commands as well as manage the data packet transfers from the producers to the consumers.

As shown in Figure 2, the BCU subsystem consists of three main functional blocks: processing, DMA, and I/O. The processing section handles all the commands coming from the SC. The processor is responsible for functions such as error detection, special-purpose data transfers, control and status transfers, and communication with the system controller. The BCU's main task is to perform high-speed data block transfers. To accomplish this, it needs a DMA to meet the speed requirements. When the BCU is interrupted, the DMA gains control of the bus, picks up information from various producers, and latches it onto the consumer bus at maximum speed.

Producer interface unit

The producer interface interconnects the producing devices and the distribution network. The data generated by individual producers is collected by a producer interface unit, or PIU, which is capable of simultaneously servicing all producers connected to it. Each PIU must provide the appropriate interface to match the characteristics of each producing device, assemble data in a special format for the BCU, accept information from the SC via the BCU, and send it to its producing devices.

Each PIU can service many producers having different input/output characteristics. In the current design, the PIU supports the following interfaces: one multipurpose parallel interface, two RS-232C type serial interfaces, and one special-purpose interface that is accommodated via a reserved wire-wrap section. The multipurpose parallel interface has a flexible structure; a number of programmable control lines to synchronize data transfers are provided. Also, data transfers can be supported either by an onboard FIFO or directly by the PIU's processor.

The main component of the PIU is a true dual-port read/write memory. This memory can be accessed simultaneously by both the BCU and the PIU and is used for transferring control and data information between them. Since multiple PIUs could be connected to the common bus, bidirectional tri-state buffers are provided on the outputs of the dual-port memory. In the event that both the PIU and the BCU attempt to access the same control memory location simultaneously, the BCU receives priority, and the PIU access is delayed until the BCU has completed its operation.

Consumer interface unit

The BCU transfers information packets collected by the PIU at high speed to the consumers. The in-

formation goes to every consumer. Each consumer has a consumer driver board connected directly onto the consumer port (Figures 1 and 2). Each driver board sends the information via twisted-pair cables (differential voltages) to the consumer receiver board (which isolates the signals using optocouplers). Data is transferred to the consumer interface unit (which provides the intelligence to sort the received information). The large distances involved and the large number of consumer devices (up to 10) forced us to minimize the number of control lines and to select a communication protocol that was independent of the consumers. The communication protocol between the BCU and the consumers permits continual availability of the data on the bus, together with its clock. The BCU asserts the block transfer start signal and a clock signal at the middle of each data word transfer.

Any interested consumer may access this data; however, a consumer cannot interrupt the transfer. So to verify consumer status, we provided three status lines, as well as one request line, for the consumer to indicate an exception condition. Whereas each consumer has a dedicated request line, we kept the three status lines common to all consumers. If one or more request lines are asserted, the BCU selects the appropriate consumer request condition by enabling its status buffer.

Data block transfers

To transfer large blocks of data from different producers, we built a distribution network with a variable memory organization and shared banks of data between the BCU and the PIUs. Many special features were incorporated into the design to facilitate the address allocation to each data word: bank switching, shared addressing, programmable selection of the PIUs, and the DMA "latch while reading" option.

In general, each PIU collects data from its producers. When a PIU is interrupted by a particular producer, the producer's data is immediately stored at its allocated address, inside the data block *x*.a (*x* is the PIU number) of the dual-port RAM (shared by the BCU and the PIU, see Figure 3). Once the entire group of data that this particular PIU has to collect fills the allocated addresses inside the data block, the PIU switches the data block flag bit. Next the PIU begins writing into data block *x*.b, allowing the BCU to read block *x*.a. By reading (BCU) and writing (PIU) different blocks, the system has no need for a complex control logic. The PIU works at maximum speed, not interfering with the BCU. The BCU reads in the latest available complete data block, always knowing that the data is valid. The PIU sets up a data block flag bit to indicate when the data has been updated.

After the data has been sorted by the PIUs, the BCU can initiate data block transfers. This is done on a regular basis (2 KHz). When the BCU receives a start signal, the DMA takes control of the bus, by supplying addresses and latching data words directly into the consumer port during the read cycle. Thus the DMA allows the bus to work at twice the speed (the current speed is 16M bits/s, but we expect it to double in the final system).

The data sent to consumers contains two types of information: control words like the block type, status, and information coming from the BCU memory, and the data block itself, collected from various PIUs. The data block size is completely programmable; it contains a variable number of control words (BCU memory) plus a variable number of data words (PIU dual-port RAM). Because the BCU memory is adjacent to the PIU data, the DMA easily transfers the BCU data, and afterwards, the PIU data. We added a special feature to the BCU to increase the flexibility of the block transfers: The PIU selects data on a word basis, not on a block basis. Therefore, a data block is a selection of independent words, each one obtained from any PIU. The block, as a whole, becomes a group of PIU data packets (seen in Figure 3). To make this possible, every BCU data block configuration must provide its own sequence of PIU selection. Up to 14 block configurations can be programmed, with immediate access to two of them by the DMA (odd blocks on the first channel, even ones on the second) and to the others by changing the starting address inside the DMA.

Finally, when the DMA has put the data on the bus, and has latched it onto the consumer port, the consumer driver board amplifies the signals and drives them to the consumer interface unit, 200 feet away, at 16M bits/s (projected speed of 32M bits/s). The CIU, after receiving the block, decodes the block type word and the status word (status of PIU data) and decides whether it should keep the block or not. If the test is positive, the data can be distributed to its allocated consumer, at the consumer's speed.

Figure 3 shows the configuration of the dual-port memories of each PIU. The design incorporates two levels of memory selection. The PIU controls the first level and forces the BCU to read a predetermined PIU data block (*x*.a or *x*.b, *x* being the PIU number) while the PIU is writing into the other (*x*.b or *x*.a). The BCU sees only one PIU data block, but the PIU selects the one (a or b) to be read by the BCU. The second level of memory selection is located on the BCU side. The PIU-selected logic specific to each BCU data block contains the PIU number associated with each data word inside the BCU data block. The resulting BCU data block constitutes an arrangement of words coming from different PIUs according to the predetermined pattern of selection. This method permits up to 14 BCU data blocks, seven of which are composed of PIU data block type one and seven, of block type two.

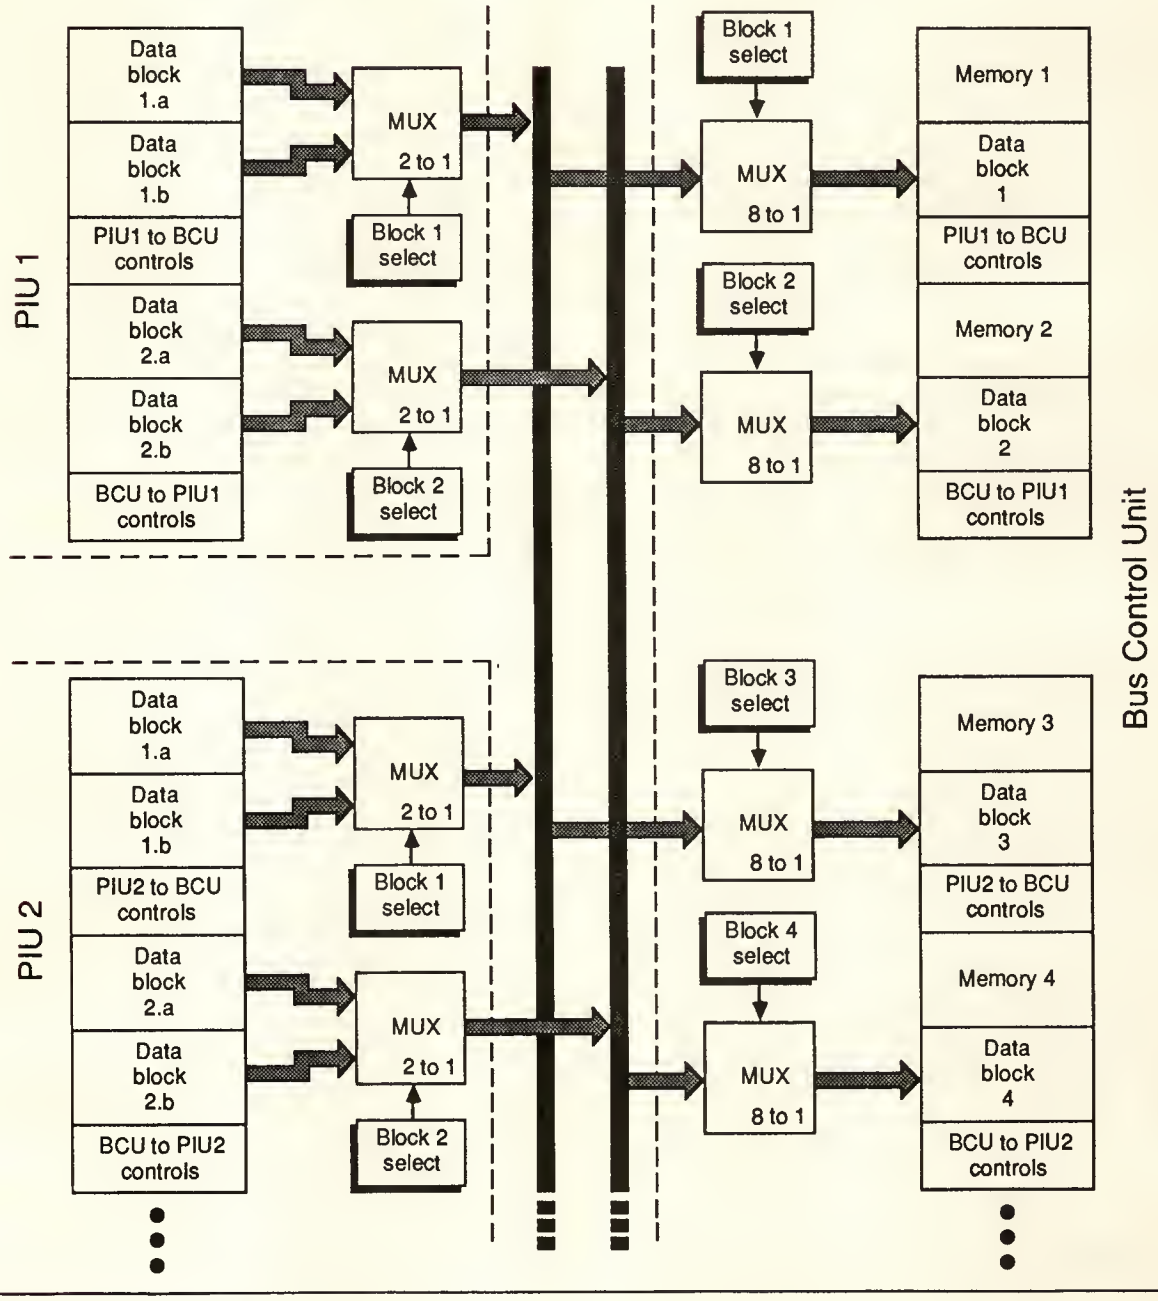


Figure 3. PIU/BCU shared memory.

BCU software: command execution

When the BCU is not interrupted by the data block transfers, it executes the commands requested by the system controller (see Figure 4). One register controls every transfer during the execution of a command: the *control* register defines when the operation level is active, when the system controller interface is active, and when the BCU transmits to the destination all of the information received from the source. One feature of the command execution software is that it can keep track of special requests while the system controller interface is busy. It permits the BCU to detect system errors and interrupt the command execution.

A transaction involving the system controller (which might be the source or the destination) requires two steps. First it obtains information from

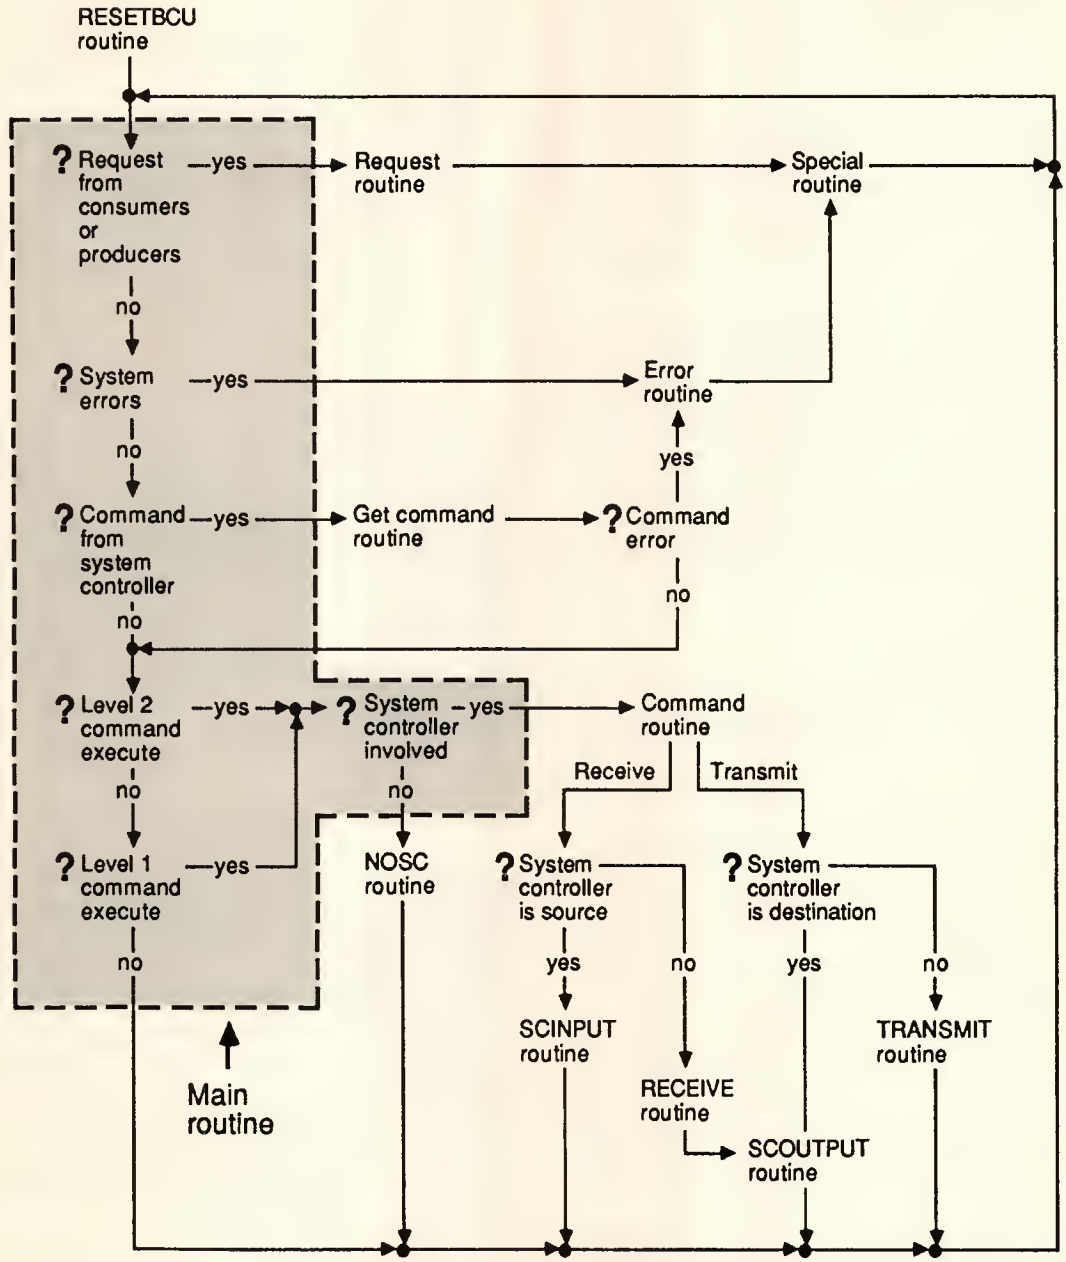


Figure 4. BCU software flowchart.

the source and stores it in the BCU memory. Next it sends this information to the destination. After receiving the command word, the control register activates the current operation level. The command word is decoded by the BCU. If the SC is not involved in the operation, the NOSC routine completes the execution of the command. Otherwise, the COMMAND routine initiates the first half of the command: the reception of the control words from the source to the BCU memory. The SCINPUT routine (SC is the source) or the RECEIVE routine (SC is the destination) receives the control words. When reception is finished, the COMMAND routine is called again by the MAIN routine to finish the command execution with either the SCOUTPUT or the TRANSMIT routine.

We have described a distribution network that facilitates reliable, high-speed data transfers for systems with multiple sources and destinations. To support this goal, we developed an elegant system bus structure, which we call a double bus star interconnection. Such a structure provides many benefits derived from both the single bus as well as the star type of routing mechanism. It shows great promise for expansion to multibus star structures for other applications that could be characterized by the producer/consumer model described here.

We intend to continue improving the system interconnection structure as technology develops. The system bandwidth is directly related to these improvements, as is the speed of operation of the bus control unit. The BCU, at present, handles all the DMA control, but in later improvements the other modules on the buses could conceivably assume the DMA function. In addition, we also plan to consider increased data rates, data quantity, and data transmission distances.

Acknowledgments

We thank Gordon Marwood, Ross Fines, and Denis Lamothe for their contributions to this project.

References

1. E.T. Fathi and M. Krieger, "Multiple Microprocessor Systems: What, Why and When," *Computer*, Mar. 1983, pp. 23-32.
2. G.A. Anderson and E.D. Jensen, "Computer Interconnection Structures: Taxonomy, Characteristics, and Examples," *Computing Surveys,* Dec. 1975, pp. 197-214.
3. K.J. Thurber et al., "A Systematic Approach to the Design of Digital Bussing Structures," *AFIPS Conf. Proc.,* 1972, Montvale, N. J., pp. 719-740.
4. D.R. Powell, "Dependability Evaluation of Communication Support System for Local Area Distributed Computing," *Proc. 12th Ann. Int'l Symp. Fault Tolerant Computers,* IEEE, New York, June 1982.
5. C. Weitzman, "Distributed Micro/Minicomputer Systems, Structure, Implementation and Applications," Prentice-Hall, Englewood Cliffs, N.J., 1980.
6. E.T. Fathi and N.R. Fines, "Real-Time Data Acquisition, Processing, and Distribution for Radar Applications," Real-Time Symp., Computer Society of the IEEE, Washington, D.C., Dec. 1984.



Eli T. Fathi is president of Applied Silicon Inc. Canada. He is currently working on the development of application-specific integrated circuits for diagnostics and testability. His main research interests are in the areas of advanced computer architectures and high-speed DSP modules.

Fathi received his BASc and MASc degrees in electrical engineering from the University of Ottawa in 1978 and 1981. He is a member of the Association of Professional Engineers of Ontario and of the IEEE.



Eloi Bosse is a research engineer for the Communications Research Centre, where he is working on developing leading-edge algorithms for radar signal processing. High-performance distributed computer architectures for real-time applications are his main research interests.

Bosse received a BScA and MSc in electrical engineering from the Universite Laval de Quebec in 1980 and 1981. He is a member of the Ordre des Ingenieurs du Quebec.



Jean Caseault is a senior design engineer for Applied Silicon Inc. Canada. When he worked on the project reported here, he was a development engineer for the Communications Research Centre in Ottawa, Canada. He has worked on a number of sophisticated microprocessor-based systems for an experimental radar. His interests are in the areas of distributed computer architectures, special-purpose controllers, and VLSI chip development for DSP applications.

Caseault received a BScA in electrical engineering from the Universite Laval de Quebec in 1982 and is a member of the Ordre des Ingenieurs du Quebec.

Questions regarding this article can be directed to Eli T. Fathi, Applied Silicon Inc. Canada, 310-2255 Ft. Laurent Blvd., Ottawa K1G 4K3, Canada.

Reader Interest Survey

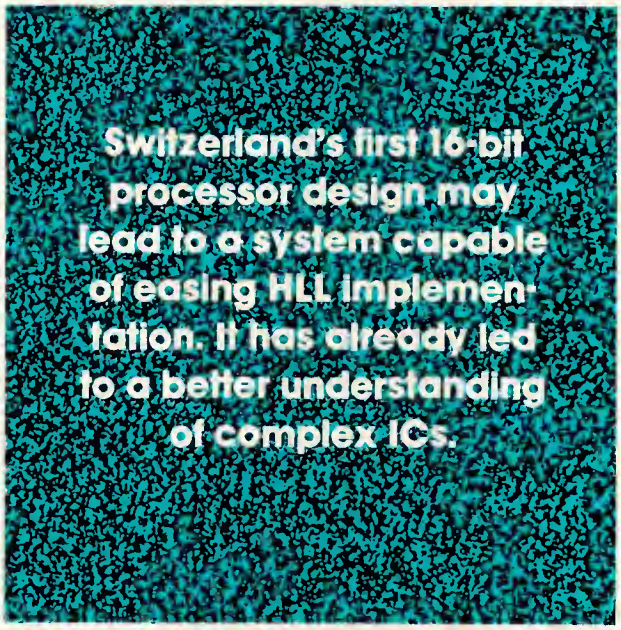
Indicate your interest in this article by circling the appropriate number on the Reader Interest Card.

Low 162 Medium 163 High 164

FEATURE

A General Heap Processor

Eduardo Sanchez, Patrick Sommer, Jacques Menu, and Christian Iseli
Ecole polytechnique federale de Lausanne



At the Ecole polytechnique federale de Lausanne in Switzerland we have recently completed the design of a 16-bit processor in two versions, bit-slice and VLSI. The goals determined at the beginning of the project were to:

- acquire a new capability (no logic design project of such a large scale had ever been undertaken in the school);
- develop design tools (microassemblers, simulators, graphic editors) that would be usable for teaching and research in the future;
- encourage software and hardware teams to cooperate in finding principles common to languages, to architectures, and to design methodologies;
- test, on a complex circuit, the integrated circuit design methods that the school developed in cooperation with the Centre suisse d'electronique et de microtechnique de Neuchatel (CSEM, Switzerland);[1] and
- obtain as a final product a processor that would ease implementation of high-level languages and particularly of the Newton language developed at the school.[2]

We chose to facilitate the implementation of the following high-level-language characteristics:

- dynamic data structures management;
- coroutines and procedures management, namely the allocation of their activation blocks;
- detection and handling of runtime errors (overflow, use of uninitialized variables, nil pointer dereferencing); and
- debugging facilities.

These choices led to a quite classical, microprogrammed processor architecture. However, we organized data memory management in the form of a *heap,* which is completely ruled by dedicated instruction.

Memory management

Most of the actual high-level languages are block-structured languages.[3] A program is decomposed into blocks, called procedures or routines (within which local variables can be declared) known only in their environment.

This property, present in Pascal for example, has consequences for the physical implementation of the language. A memory allocation for a local variable is necessary only when its procedure (the block which contains it) is active. However, this economy of memory leads to a problem: The zone of the memory that contains the local variables cannot be statically allocated. The classical solution to the problem is the implementation of the data memory in the form of a stack.[4] In that way, to each call of a procedure, designers allocate a block of memory placed at the top of the stack (the activation block) for its local variables (the stack increases). When the procedure ends its execution, the memory block becomes free (the stack decreases) and can be used again by another procedure. The memory's location is then used in the optimal way.

When one is confronted with dynamic variables (managed by the New and Dispose procedures of Pascal), management of the memory is a little more complex because the variables can live longer than (or can survive to) the procedure that has created them. The management of the coroutines leads to the same problems.

0272-1732/87/1000-0029$01.00 © 1987 IEEE

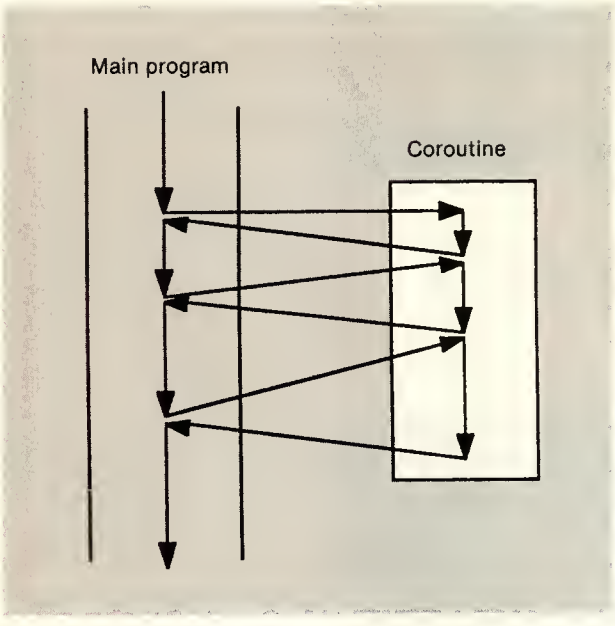


Figure 1. Execution of a coroutine.

The essential difference between a procedure and a coroutine is the following. When a procedure is called, it is executed from beginning to end; that is to say, with only one entry point and one output point. However, a coroutine can be stopped at any point and control goes back to the routine that has called it. At the next call, the coroutine comes back to the point from which it left, keeping its previous state (Figure 1). The stack is no longer a satisfactory solution for the problem of memory allocation for variables, since the memory blocks are no longer necessarily released in the inverted order of their allocation.

As a result we resorted to a heap form of managing memory. Blocks of the memory are allocated at each call of the procedures or of the coroutines and at each creation of dynamic variables. The blocks of the heap are sequentially allocated from one extremity of the memory, and the heap increases in the direction of the other extremity of the memory. As soon as one has finished using a block of the heap, one marks it as being free. The heap then composes a suite of free and occupied blocks, in an aleatory order. As the blocks are sequentially allocated, a moment arrives when the heap occupies the whole memory. One then needs a garbage collector that compacts the blocks occupied on one side of the memory and the free blocks on the other. The allocation of new blocks begins again at the last occupied block.

The notion of general heap leads to only one philosophy of memory management, although one generally finds a stack and a heap in the usual implementations of high-level languages.

Organization. The entire data memory available to the processor, lying between the *heaplow* and *heaplim* addresses, is organized in the form of a general heap: memory blocks are allocated through a basic instruction named Allocate.

As seen in Figure 2, the blocks are linked together, forming two types of chains:

- *The process chain*. The first block is the system block and the last is the current process block. Between these two blocks are all the blocks of the processes that have been suspended but can be reactivated by the program.
- *The procedure chain*. Each process has its own procedure chain, which is organized similarly to the stack of the usual high-level-language implementation.[4]

It is guaranteed that any data created at some place in a program will remain alive as long as it is needed by the execution of that program. In particular, there is no means to destroy data explicitly (compare Dispose in Pascal), and dangling references to destroyed data cannot exist. When a block is no longer useful, the space it occupies is automatically recovered at the next garbage collection/compaction.

Any reference at a memory element is made through a double pointer named *heep*. The first word of a heep is named the address part and contains the initial address of a block in the heap. The second word, which is named the complement part, contains additional information, generally as an offset inside the block.

The state of the program is held at every moment in four heep registers:

- *System*. The address part contains the address of the system block; the complement contains nil.
- *Global*. The address part contains the address of the program block; the complement contains the process type.
- *Current*. The address part contains the address of the active process block; the complement contains the process number.
- *Local*. The address part contains the address of the active procedure block; the complement contains nil.

The structure of a block. The base element of the heap is the data block. A block is requested through the Allocate instruction when a quasiparallel process (coroutine) is created, when a procedure is called, and at each dynamic data allocation (character strings, unbounded sets, rows, queues, associative tables, and so on).

Figure 3 shows the structure of blocks in memory. Each block is composed of:

- a four-word header;
- a part containing the heeps, that is, references to other blocks; and

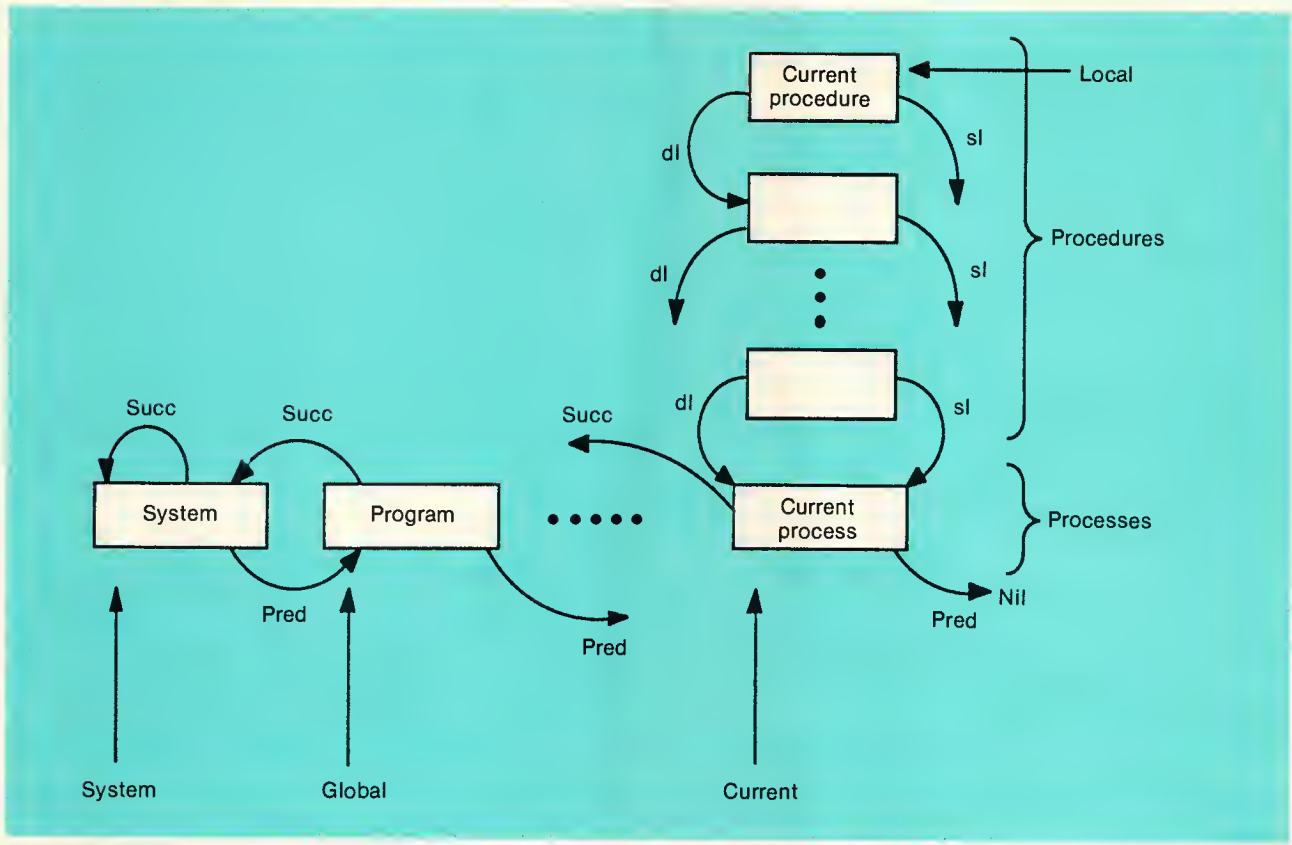


Figure 2. Block linkage.

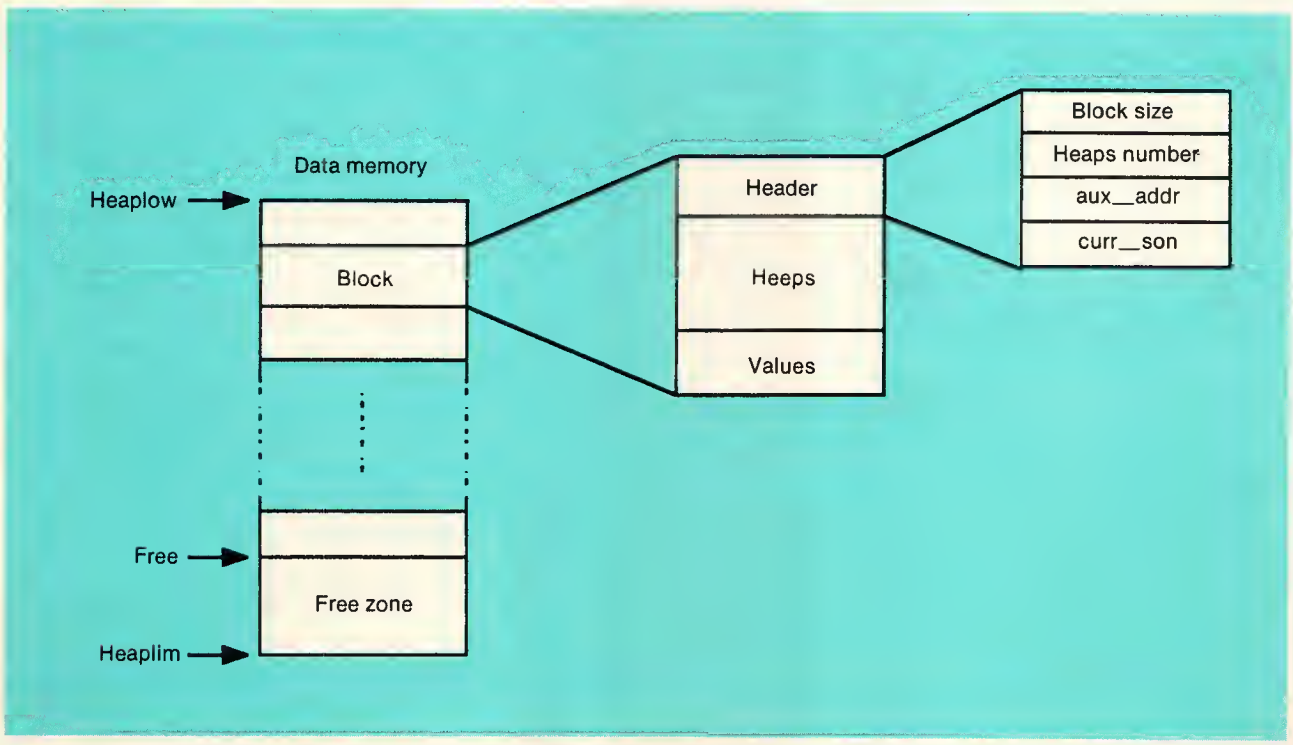


Figure 3. Memory organization.

• a self-contained data part, which groups data not referring to any other block in the heap, such as characters, Booleans, numbers, and the like.

The header does not contain program data, only information necessary to memory management. The first word gives the size of the block (in number of words), and the second word, the number of heeps. The two last words are temporary variables used during the garbage collection.

The block of a process has at least five heeps used to link this process to the process chain in the heap:

• *Static link*. The address part points to the block of the process or procedure where this process was declared; the programmer is free to use the complement part.
• *Dynamic link*. The address part points to the block of the procedure that was active at the time of the suspension of the process or calling routine; the complement part contains the return address in the code.
• *Successor link*. The address part points to the block of the process that will be executed after the current process is suspended; the complement part contains the process number.
• *Predecessor link*. The address part points to the block of the process that was executing before the current process; the complement part contains the process number.
• *Global link*. The address part points to the block of the program that owns the process; the complement part contains the process status.

Only two heeps are necessary to link the procedure's blocks, the static link, and the dynamic link.

The system block contains, in its self-contained data part, eight useful values for the debugging of programs. They are:

• free_sys_save, the address of the top of the heap (free) at the time the system launched the program. At every instant free gives the address at which the next block can be allocated on the heap;
• free_prog_save, the value of free at the time the program stopped;
• pc_save, the value of the program counter at the time the program stopped;
• ms_save, the processor status at the time the program stopped;
• current_save, the value of current at the time the program stopped;
• global_save, the value of global at the time the program stopped;
• local_save, the value of local at the time the program stopped; and
• heaplim_save, the value of heaplim, that is, the address of the top of the data memory. The size of the data memory is heaplim – heaplow.

Garbage collection. The heap may be considered as a directed graph whose nodes are blocks and whose arcs are heeps. The root block is always reachable via Local, one of the processor's heep registers. A block is said to be reachable if the graph contains a path from Local to the block. Note that circular passes can occur in the heap and that the memory manager must cope with them.

The only way for a program to indicate that a block is no longer needed is to cut an arc leading to it, for example, by assigning a distinct value (nil) to a heep that pointed to it. However such a block becomes unreachable only when the given arc is the only one leading to it. By cutting such an arc, it is also possible to make a whole subgraph unreachable, if the only references to its nodes originated from the arc just cut.

The physical data memory is divided into two logical parts (Figure 3); at one end is the heap zone, between heaplow and free, containing contiguously all the blocks currently allocated, whether reachable or not. The other end, between free and heaplim, is a free zone from which space can be fetched to allocate a block whenever requested. The heap zone grows as long as block requests can be satisfied by the free zone, moving the logical barrier toward the free end of memory.

When the free zone becomes too small to satisfy a request, the process of garbage collection begins. It successively executes the following phases:

1) Determine which blocks in the heap zone are reachable, starting from Local.
2) Compute the future address of every reachable block, that is, the address that will be assigned to it after compaction.
3) Update the heeps, replacing the current address field by the new address computed in the previous phase.
4) Move all reachable blocks to the heap end of memory, thus creating a single, new free zone.

If the block request that caused garbage collection to occur can be satisfied, the execution of the program proceeds normally. Otherwise the heap overflow is flagged and the program aborts.

The first phase (marking reachable blocks) is the only one that needs to go through the heap, starting from Local and following the arcs; the last three are merely linear scans through the heap zone. This process of going through the heap is conceptually recursive, since each heep in a block just marked points to a block that has to be processed too.

How can we handle the (implicit or explicit) stack needed to implement this recursive algorithm when the physical memory is precisely full? We do this by reserving enough space in each block to contain the origin of the arc that led to it. This reversed arc is stored in the third word of the header, named

We built the Hip1 prototype of our processor, using bit-slice circuits, around the Am29203.

aux_addr. The word is set the first time the marking phase visits a reachable node. Since this word is normally nil, a separate Boolean field to mark reachable blocks is not necessary. The last three phases recognize the blocks as having a nonnil value in aux_addr.

How in turn do we go through the subgraphs originating from a node? This we do with a For loop whose index is stored in the fourth word of the header, named curr_son.

Multiple arcs leading to the same block and circular paths in the graph are easily handled by not processing heeps contained in an already marked block.

This heap management can be seen as a generalization of the Lisp 2 algorithm described by Knuth,[5] whose performance Cohen and Nicolau have assessed.[6]

Instructions for memory management. The base instruction in memory management is Allocate *m*,*n*, which allocates a block of *m* words and which contains *n* heeps. This instruction causes garbage collection if necessary.

We provide an instruction Garbage as a convenience, so programmers can launch a garbage collection at any time.

Process management instructions. The first instruction to be executed by the processor (INIT_RUN) initializes the blocks of the operating system. To launch a user program, the system executes the instruction ENTER_PROG adr, where adr is the starting address of the program code. This instruction also begins garbage collection, to furnish the maximum free data memory to the program.

The program code, possibly generated by a compiler, must first allocate the block for the program with an Allocate instruction and then use the instruction ENTER_SECT adr to start the execution of the program at the address adr. The latter instruction also updates the processor registers and the links between the program and the system.

At the end of the program the instruction END_PROG returns control to the system.

The program may be divided into processes (or coroutines) and into procedures. The instruction ENTER_PRSS adr starts the execution of a process at the address adr and updates the processor registers and the links of the process. The process block must have been created before that by the Allocate instruction. The process execution ends with the instruction Terminate, which then activates the successor of the process.

We reactivate a stopped process with the following four instructions:

- Activate pr. The process containing this instruction will become the successor of the activated process. The pr parameter contains the address of the activated process block and the address of the code to be executed.
- Resume pr. The process containing this instruction is detached. The pr parameter contains the address of the block of the process that will be activated along with the address of the code to be executed.
- Return. Control passes from the current process to its successor. The current process is detached.
- Exchange. The current process is exchanged with its successor.

Procedure management instructions. Two instructions can be used to activate a procedure: ENTER_PROC and ENTER_NAME. The former takes the starting address of the code of the procedure as a parameter, and the latter takes the address of the block of the procedure and the starting address of the code as a heep parameter. In both cases the block of the procedure must formerly have been allocated with an Allocate instruction.

The instruction EXIT_PROC terminates the execution of a procedure.

Bit-slice implementation

We built a prototype of the processor, using bit-slice circuits, around the ALU Am29203; its name is Hip1.[7] The Harvard-like processor has two zones of memory that are physically separated:

- the program memory, which has a size of 64K words of 16 bits; and
- the data memory (the heap), which has a size of 64K − 1 words of 16 bits. The entire 64K words are not used because the address 0 is used as the value for nil (that is to say, that heaplow = 1 and heaplim = $2^{16} - 1$).

The two memories are word addressable.

Input and output are treated separately: in this case, only the low-order byte of the address and data buses is used, allowing 256 input units and 256 output units of 8 bits each.

The microprogrammed control unit uses the Am2910 sequencer. The microcode is very horizontal, allowing optimal use of parallelism; its size is 1K × 72 bits, of which 116 microinstructions are used to implement the garbage-collection algorithm.

System management
 The INIT_RUN instruction allocates and initializes the block for the operating system. It must be the first instruction to be executed when the processor is powered up.
Program management
Process management
Procedure management
Memory management
Register transfer
Loading a register from the data memory
Loading a register from the program memory
Storing a register in the data memory
Storing a register in the program memory
Register intialization
Integer arithmetic operations
Cardinal operations
Byte operations
Bit operations
Shift and rotate operations
Logical operations
Jump operations
Input/output operations
Miscellaneous operations
 In addition to the NOP operation, there are two debugging instructions:
 DUMP adr updates the system block, namely the eight values described earlier, and passes control to the debugging procedure, which starts the code at the address adr.
 ERROR #n updates the system block, stops the execution of the program, and returns control to the operating system. The value *n* is assigned to the program termination status. *n* is a user-defined error code ($16 \leq n \leq 225$; the system reserves the codes 0...15).

Figure 4. Classes of instruction.

The processor is physically realized on a double VME-format board; another board implements the interface between the processor and a VMEbus, thus allowing communication between the processor and other commercial boards like memory and interface boards.

Data types. Hip1 handles four data types: heeps, integers, bytes, and cardinals (unsigned integers). Three types of registers handle these different data:

- eight heep registers, H0 through H7. Three of these, H5, H6, and H7, are used for Global, Current, and Local addresses. A System register does not exist; the systems block always starts at address 1 (heaplow);
- 16 word registers, W0 through W15; and
- 16 byte registers, B0 through B15.

Addressing capabilities. The data may be addressed in three different ways with Hip1:

- *Indirect addressing*:

 address = Hi.a + offset

where Hi.a is the address part of the heep Hi, and offset is a 16-bit value given by the instruction. An example is

```
LOAD_INDIR      LOCAL,#12,W0
M[LOCAL.a + 12] → W0
STORE_INDIR     H4,LOCAL,#8
H4 → M[LOCAL.a  + 8]
```

- *Referenced addressing*:

 address = Hi.a + Hi.c

where Hi.c is the complement part of register Hi. Example:

```
LOAD_REFER      H4,W0
M[H4.a + H4.c] → W0
STORE_REFER     B2,H3
B2 → M[H3.a + H3.c]
```

- *Indexed addressing*:

 address = Hi.a + Wj − Hi.c + 4

where Wj is a word register and 4 is the size of a block header. This addressing mode is used to access array elements. Hi.c contains the lower bound of the indices, while Wj contains the index of the particular element to be accessed. Moreover, the register Wj may be post-incremented or pre-decremented. For example:

```
LOAD_INDEX      H2,W3,W0
M[H2.a + W3 − H2.c + 4] → W0
LOAD_INDEX      H1, (W3+),B2
M[H1.a + W3 − H1.c + 4] → B2
W3 + 1 → W3
STORE_INDEX     W1,H3,( −W2)
W2 − 1 → W2
W1 → M[H3.a + W2  − H3.c + 4]
STORE_INDEX     H0,H2,W4
H0 → M[H2.a + 2(W4 + H2.c) + 4]
```

Of course, direct addressing is not available to reference a data block. The only block that has a fixed place and that cannot be moved by garbage collection is the system block. But direct addressing does exist to read constants or to store code in the program memory.

Instruction set. The 136 instructions of Hip1 may be classified into 20 groups, as seen in Figure 4. Eight execution errors detected by the instructions' microcode are:

- OVERFLOW, for integer arithmetic operations;
- DIVZERO, division by zero;
- ERRBORNE, out-of-bounds error in a JUMP_INDEX instruction;
- VALIND, use of an uninitialized value (nil value);
- ADRIND, use of an uninitialized address (nil pointer);

• NODETACH, trying to activate a process that was not detached;
• ERRETURN, trying to return to the system process; and
• ERREX, trying to exchange with the system process.

If one of these eight error conditions is detected, the microcode stops the execution of the program, gives the control back to the system, and updates the program termination status and the system block. Program debugging is thus facilitated.

The instructions that modify the content of any register update the processor flags, allowing conditional jumps. The processor flags are:

• OVR, overflow;
• CARRY;
• NEG, sign bit;
• EQUAL, zero bit;
• LESST, less than bit; and
• GREAT, greater than bit.

Development tools. The very particular characteristics of a microprogrammed system ask for specialized tools to help the design and the debugging of such a system.[8] We wrote and tested the microprogram for the Hip1 processor thanks to the two following tools, which were designed by our staff:

• *A variable-definition microassembler.* In a first pass, the microassembler receives the microinstruction format definitions (widths, field places, default values) and the interpretation of symbolic names (opcodes, definitions for the microoperations and parameters). In the second pass it receives the microprogram written using the language defined during the first pass and outputs its translation in object code.

• *A simulator.* The simulator takes the assembled microcode and a description of the processor at the register-transfer level (list of the registers and description of their interconnections by a list of microoperations). It then uses this information to simulate the functional behavior of the microcode.

These tools are not dedicated and allow us to design microcode for virtually any processor. They are written in Pascal and used on VAX, HP9836, and Sun computers.

VLSI implementation

At the same time we were designing the bit-slice implementation, a part of our team worked on a VLSI implementation to practice and test new design methods for integrated circuits.[9]

Internal architecture. Basically, the VLSI implementation architecture of the processor is the same as the bit-slice implementation. But, in general, the different functional blocks have less possibilities than their bit-slice equivalent. For example, the microprogram sequencer allows only a single level of subprogram.

The microcode uses about 1300 words of 36 bits: longer and narrower than that used in the bit-slice version. This difference in size comes from the vertical characteristics of the microinstructions. The control microoperations are separated from the processing microoperations.

The execution of a microinstruction decomposes into four phases:

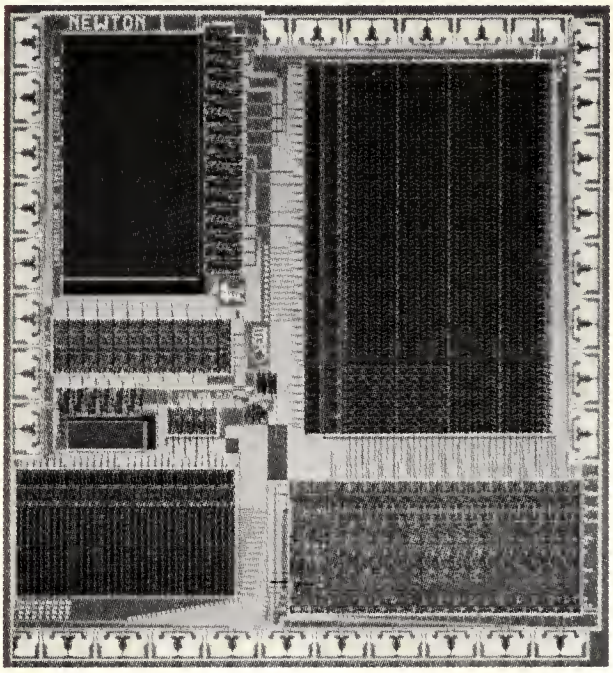
• loading of the micro-PC slave register; micro-ROM preload;
• micro-ROM reading;
• processing and loading of the data in the accumulator; and
• loading of the destination registers; loading of the micro-PC master register.

External architecture. From the user point of view, some differences exist between the two processors. The bit-slice version is a Harvard machine, with two separate memories for the data and programs; the VLSI version is a von Neumann machine with a single memory. Such a difference implies a slightly different heap management. In the VLSI version the programs are stored starting at address zero, toward the high end of memory. The first block of the heap lies at the high end of memory, and the heap grows toward the start of the memory. Except for some details, the garbage-collection algorithm stays the same.

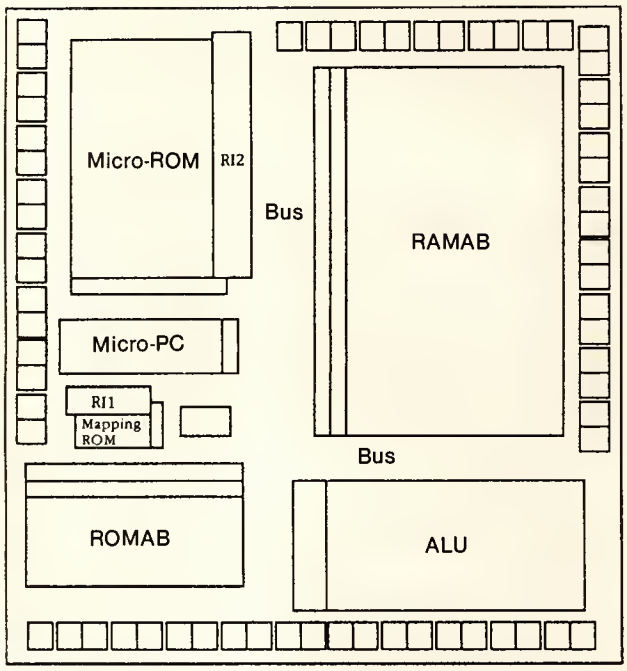
The VLSI version of the Hip processor also uses 16-bit words but has only eight registers of each type (byte, word, and heep), because of the space needed by these registers. They occupy about 20 percent of the chip surface.

Methodology. The circuit design follows a method developed at our school that avoids the need of a logical scheme.[1,10] The starting point is the Karnaugh table of each function. The sequential parts are realized asynchronously (without explicit memory elements) and may also be characterized by a Karnaugh table. By separately computing the 1 and 0 covering of the table, we obtain the network equation of the p and n transistors of a CMOS circuit. The covering of a Karnaugh table having up to 10 variables may be computed by a program.

We draw the layout on a symbolic grid consisting of a horizontal conductor (metal) and a vertical conductor (polysilicon). At the intersections may lie a crossing, a contact, or a transistor that has its gate bound to the vertical conductor. A program automatically translates this symbolic drawing into a



(a)



(b)

Figure 5. Chip microphotograph (a) and floor plan (b).

geometrical layout. Since the program takes care of the layout rules, designers do not need to worry about them. They can concentrate on placing the variables, and in doing that, take into account the connections to the neighboring cells.

This methodology allows a very short design time. The design and drawing effort of this processor may be estimated to be about eight man-months. The chip is implemented in 3-micron SACMOS technology (self-aligned CMOS).

Chip floor plan. Figure 5 shows the floor plan and the microphotograph of the chip. The block named ALU contains the entire processing part, except the register bank (RAMAB). The micro-PC block contains the control part, except the mapping-ROM and the micro-ROM. RI1 and RI2 are registers (12 and 36 bits) placed at the ROM output for testing purposes. They normally operate as parallel input-output registers but can be turned into serial mode. This capability allows us to examine or modify their contents. The ROMAB block contains constants (same bus address as RAMAB).

Here, we give some technical data to characterize the chip. It contains 71,551 transistors of which 46,944 are assigned to the micro-ROM, 14,576 to the register bank, and 3218 to the ALU. The chip surface occupies 20.1 mm^2. The bus and the contact pad take 8.72 mm^2 of this surface, the register bank, 4.65; the micro-ROM, 1.97, and the ALU, 1.94.

The goals set at the start of the project were wholly fulfilled.

1) For the first time in Switzerland, an integrated circuit for a 16-bit processor has been designed and realized.

2) A new course, named "Conception des processeurs" (processor design), benefits from this cooperation among software and hardware people. Students become acquainted with a global vision of computer architecture and its relations with high-level languages.

3) The compiler for the Newton language is not yet completed, and it is too early to give benchmark results. But the first tests of the processor in its bit-slice version are encouraging. As a matter of fact, a cross-compiler for the Pascal-S language,[11] a subset of Pascal, has been written on a Sun workstation, and the famous sieve of Eratosthenes can be executed in 6.3 seconds with a 5-MHz clock. (This program appears in the adjoining box.) The same program, compiled with Turbo Pascal, takes 5.1 seconds on a Macintosh Plus computer, which has an 8-MHz clock.

We have already started realizing the 32-bit versions of the processor, bit-slice as well as VLSI, and we foresee the first results early in 1988. We have also started the design of a processor adapted to the Prolog language, taking advantage of the same methodology and of our experience in this project.

The Sieve of Eratosthenes Program for Computing Prime Numbers

Here is the famous sieve of Eratosthenes program to compute prime numbers.

```
program primes;

const
  size = 8190;

var
  flags : array [0 .. size] of Boolean;
  i, prime, k, count, iter: integer;

begin
  writeln('10 Iterations');
  writeln('Hit return to start...');
  readln;
  for iter := 1 to 10 do begin
    count := 0;
    for i := 0 to size do
      flags[i] := true;
    for i := 0 to size do
      if flags[i] then begin
        prime := i + i + 3;
        k := i + prime;
        while k <= size do begin
          flags[k] := false;
          k := k + prime
        end;
        count := count + 1
      end
  end;
  writeln('There are ', count, ' primes.')
end.
```

And now, here is the result after grinding the above program with the Pascal-S compiler. The produced code has been commented, and blank lines have been inserted to make it clearer.

```
                ALLOCATE        #8204, #5               ; Program block allocation
                ENTER_SECT      primes$1                ; Go start the program

true$0          =               -1                      ; Program constants
size$1          =               8190
i$1             =               4+10+0                  ; Offsets of the program vars
prime$1         =               4+10+1
k$1             =               4+10+2
count$1         =               4+10+3
iter$1          =               4+10+4
flags$1         =               4+10+5


                .ENTRY          primes$1                ; Main program starts here

                ALLOCATE        #8, #2                  ; Procedure call to the
                ENTER_PROC      init_pascal_io          ; standard I/O initialization

                ALLOCATE        #13, #2                 ; writeln('10 iterations')
                LOAD_ADDRESS    s$0$strg, W15
                STORE_INDIR     W15, H0, #8
```

```
                ENTER_PROC      write_string
                ENTER_PROC      writeln

                ALLOCATE        #13, #2             ; writeln('Hit return to start...')

                LOAD_ADDRESS    s$1$strg, W15
                STORE_INDIR     W15, H0, #8
                ENTER_PROC      write_string
                ENTER_PROC      writeln

                ALLOCATE        #8, #2                      ; readln
                ENTER_PROC      readln

                SET_ONE         W15                         ; for iter := 1 to 10 do (1)
                STORE_INDIR     W15, GLOBAL, #iter$1
10:             -
                LOAD_INDIR      GLOBAL, #iter$1, W15
                COMPARE         W15, #10
                JUMP_GREAT      11

                SET_ZERO        W14                         ; count := 0
                STORE_INDIR     W14, GLOBAL, #count$1

                STORE_INDIR     W14, GLOBAL, #i$1           ; for i := 0 to size do (2)
12:             -
                LOAD_INDIR      GLOBAL, #i$1, W15
                COMPARE         W15, #8190
                JUMP_GREAT      13

                LOAD_DIRECT     #true$0, W14                ; flags[i] := true
                SET_REFER       GLOBAL, #4-flags$1, H4
                STORE_INDEX     W14, H4, W15

                INT_SUCC        W15, W15                    ; end of loop (2)
                STORE_INDIR     W15, GLOBAL, #i$1
                JUMP            12

13:             -
                SET_ZERO        W15                         ; for i := 0 to size do (3)
                STORE_INDIR     W15, GLOBAL, #i$1
14:             -
                LOAD_INDIR      GLOBAL, #i$1, W15
                COMPARE         W15, #8190
                JUMP_GREAT      15

                SET_REFER       GLOBAL, #4-flags$1, H4      ; if flags[i] then
                LOAD_INDEX      H4, W15, W14
                BIT             W14, #0
                JUMP_NOT_LESST  16

                INT_ADD         W15, W15, W14               ; prime := i + i + 3
                LOAD_DIRECT     #3, W13
                INT_ADD         W13, W14, W12
                STORE_INDIR     W12, GLOBAL, #prime$1

                INT_ADD         W15, W12, W14               ; k := i + prime
                STORE_INDIR     W14, GLOBAL, #k$1

17:             -                                           ; while k <= size do (4)
                LOAD_INDIR      GLOBAL, #k$1, W15
                COMPARE         W15, #8190
                JUMP_GREAT      18
```

```
                SET_ZERO        W14                         ; flags[k] := false
                SET_REFER       GLOBAL, #4-flags$1, H4
                STORE_INDEX     W14, H4, W15

                LOAD_INDIR      GLOBAL, #prime$1, W13       ; k := k + prime
                INT_ADD         W15, W13, W12

                STORE_INDIR     W12, GLOBAL, #k$1

                JUMP            17                          ; end of loop (4)
18:             -

                LOAD_INDIR      GLOBAL, #count$1, W15       ; count := count + 1
                INT_SUCC        W15, W14
                STORE_INDIR     W14, GLOBAL, #count$1

16:             -                                           ; end of if statement

                LOAD_INDIR      GLOBAL, #i$1, W15           ; end of loop (3)
                INT_SUCC        W15, W15
                STORE_INDIR     W15, GLOBAL, #i$1
                JUMP            14
15:             -

                LOAD_INDIR      GLOBAL, #iter$1, W15        ; end of loop (1)
                INT_SUCC        W15, W15
                STORE_INDIR     W15, GLOBAL, #iter$1
                JUMP            10
11:             -

                ALLOCATE        #13, #2                     ; write('There are ')
                LOAD_ADDRESS    s$2$strg, W15
                STORE_INDIR     W15, H0, #8
                ENTER_PROC      write_string

                LOAD_INDIR      GLOBAL, #count$1, W15       ; write(count)
                ALLOCATE        #19, #2
                STORE_INDIR     W15, H0, #8
                ENTER_PROC      write_integer

                ALLOCATE        #13, #2                     ; writeln(' primes.')
                LOAD_ADDRESS    s$3$strg, W14
                STORE_INDIR     W14, H0, #8
                ENTER_PROC      write_string
                ENTER_PROC      writeln

                END_PROG                                    ; end of program

s$3$strg: -                                                 ; Strings definitions
                .WORD                   ^A / p/, ^A /ri/, ^A /me/, ^A /s./, -
                                        0
s$2$strg: -
                .WORD                   ^A /Th/, ^A /er/, ^A /e /, ^A /ar/, -
                                        ^A /e /, 0
s$1$strg: -
                .WORD                   ^A /Hi/, ^A /t /, ^A /re/, ^A /tu/, -
                                        ^A /rn/, ^A / t/, ^A /o /, ^A /st/, -
                                        ^A /ar/, ^A /t./, ^A /../, 0
s$0$strg: -
                .WORD                   ^A /10/, ^A / I/, ^A /te/, ^A /ra/, -
                                        ^A /ti/, ^A /on/, ^A /s/@8
```

Acknowledgments

The Hip processor is the fruit of the cooperation of two laboratories, one dedicated to hardware and the other to software. The following people, staff of one of these laboratories, also participated in this project: Y. Abdoun (VME interfacing), C. Bernard (printed circuits), and S. Mourtada (software). The Centre suisse pour l'electronique et la microtechnique (CSEM, Neuchatel-Switzerland) lent us its equipment to design the final layout of the integrated circuit. The Commission pour l'encouragement de la recherche scientifique and the Fonds national suisse de la recherche scientifique (Switzerland) financed a part of the project.

References

1. C. Piguet, J. Zahnd, A. Stauffer, M. Bertarionne, "A Metal-Oriented Layout Structure for CMOS Logic," *IEEE J. Solid-State Circuits,* June 1984, pp. 425-436.

2. C. Rapin and J. Menu, "The Newton Language," *SIGPLAN Notices,* Aug. 1981, pp. 31-40.

3. J. Menu, "The General Heap, A High-Level Concept," PhD thesis 495, Ecole polytechnique federale, Lausanne, Switzerland, 1983.

4. A.S. Tanenbaum, "Implications of Structured Programming for Machine Architecture," *Comm. ACM,* Mar. 1978, pp. 237-246.

5. D.E. Knuth, *The Art of Computer Programming, Vol. 1: Fundamental Algorithms,* Addison-Wesley, Reading, Mass., 1973.

6. J. Cohen and A. Nicolau, "Comparison of Compacting Algorithms for Garbage Collection," *ACM Trans. on Programming Languages and Systems,* Oct. 1983, pp. 532-553.

7. E. Sanchez, "Processeur adapte au langage Newton: version circuit en tranches," internal report, Laboratoire de systemes logiques, Ecole polytechnique federale, Lausanne, 1986.

8. E. Sanchez and P. Sommer, "Outils d'aide a la conception des systemes microprogrammes," internal report, Laboratoire de systemes logiques, Ecole polytechnique federale, 1985.

9. P. Sommer, "Processeur adapte au langage Newton: version circuit integre," internal report, Laboratoire de systemes logiques, Ecole polytechnique federale, 1986.

10. A. Stauffer and J. Zahnd, "Methodes de synthese des systemes logiques CMOS," bulletin ASE/UCS, Zurich, Jan. 1984, pp. 18-34.

11. N. Wirth, "Pascal-S: A Subset and Its Implementation," in D.W. Barron, *Pascal—The Language and Its Implementation,* John Wiley and Sons, New York, 1981.



Eduardo Sanchez is engaged in teaching and research in the areas of logical design and computer architecture at the Ecòle polytechnique federale in Lausanne, where he has been working since 1977. His current interests include compilers, microprogramming, logic programming, and Prolog-oriented computer architecture.

Sanchez received the diploma of electrical engineer from the Universidad del Valle in Cali, Colombia, in 1975 and his PhD from the Ecole polytechnique federale in 1985.



Patrick Sommer works on an Ecole project of automatic design methods for CMOS logic circuits with ordered layout. His interests lie in VLSI design and compilers.

Sommer received the diploma of electrical engineer from the Ecole polytechnique federale in 1982.



Jacques Menu lectures at the University of Geneva and is a scientific associate at CERN (the Centre europeen de recherches nucleaires) in Geneva, Switzerland. His current interests include logic programming, specification languages, and rewrite rules as a problem-solving paradigm.

Menu received his PhD from the Ecole polytechnique federale in 1983.



Christian Iseli is a diploma candidate at the Ecole polytechnique federale and is currently working on the design of a small general-purpose microprocessor to be used in watches. His research interests lie in high-level-language oriented architecture, compilers, and operating systems.

Questions concerning this article may be directed to the authors at Ecole polytechnique federale de Lausanne, 16. ch. de Bellerive, CH-1007 Lausanne, Switzerland.

Reader Interest Survey

Indicate your interest in this article by circling the appropriate number on the Reader Interest Card.

Low 165 Medium 166 High 167

FEATURE

A Fast Integer Binary Logarithm of Large Arguments

Reinhard Maenner
Physical Institute,
University of Heidelberg

The logarithm is a function defined on the space of nonnegative real numbers ($x \in \mathcal{R}$, $x > 0$). Its value is real ($x \in \mathcal{R}$, $-\infty < x < \infty$) and can be interpreted as giving the order of magnitude of the argument. We often therefore use logarithms in which variables with a large range of values are to be handled. A typical example is the processing of acoustic signals. The human ear has a logarithmic sensitivity over a range of $1{:}10^6$. An analysis example of individual Fourier components, that is, of their time dependencies, is meaningful only in considering the logarithms of these components.

Traditionally, we have executed such computations with floating-point arithmetic. All mainframes and most minicomputers use hardware floating-point processors for this task. High-end microprocessors typically exploit floating-point coprocessors to support real arithmetic. Low-cost systems, however, still emulate floating-point operations by software. Typical examples of low-cost systems are personal computers such as the Macintosh Plus or the Atari 1040 ST. Such systems are increasingly applied to control small experiments because of their low-cost graphic capabilities and good human interfaces. Here the emulation of floating-point operations and, among others, the computation of logarithms generally degrades the system throughput by two orders of magnitude.

In some cases this performance degradation can be avoided. If the dynamic range of the parameters handled can still be represented with integer numbers, a floating-point computation is not necessary. A 32-bit microprocessor usually can handle integer numbers of up to 64 bits (32 bit * 32 bit multiply), corresponding to a dynamic range of $1{:}10^{19}$. Thirty-two-bit integers still cover a dynamic range of $1{:}10^9$.

Fortunately, the logarithm is one of the real functions that can be approximated quickly by integer and logical operations. Here, I review methods to compute such approximations and propose a new method that can be implemented easily in software only, is computed very fast, and allows users to choose an optimum regarding required table space and approximation error.

Simple and fast ($10\,\mu$), this new procedure requires no hardware and features a 10^{-6} approximation error.

Previous work

In the following we will consider only integer arguments with a reasonable number of bits. If this number is too low, for example, ≤ 16, the computation of the logarithm becomes trivial, for example, using a lookup table (as discussed later). If the number of bits is too high, say, ≥ 64, the computation has no practical importance. This is true because even high-performance processors at most handle 64-bit integers (usually results of a 32-bit * 32-bit multiply) and because larger integers cover such an immense dynamic range that cannot be used in most computations.

Lookup tables using ROMs or RAMs. Because we consider only a finite set of arguments, we can in

0272-1732/87/1000-0041$01.00 © 1987 IEEE

principle compute the logarithm of each possible argument in advance and store all results in a table. A lookup table can then be used to "compute" the logarithm extremely fast. The computation time here equals the access time of read-only memories (ROMs) or random access memories (RAMs) storing the table, which is typically 50-200 ns. This method, proposed by Brubaker and Becker,[1] can however only be used with a rather limited range of arguments. Because for n-bit integers, one table entry is required for each one of the 2^n possible arguments, the necessary ROM/RAM space grows exponentially with the argument word length. Even by providing huge tables of megabyte size (the typical main memory size of today's personal computers), the argument length would be limited to a little above 20 bits.

Linear approximation. For larger arguments, some kind of computation has to be done in real time. Note that for some special arguments this computation is trivial. The binary logarithm (ld) is defined by $x = 2^{ld(x)}$. The argument is represented in binary form by

$$x = \sum_{i=0}^{n} a_i 2^i, \tag{1}$$

where $a_i \in \{0,1\}$. If therefore only a single one of the a_i's is set, say a_k, $x = 2^k$ and $ld(x) = k$, giving an integer result. Note that in this case the result equals the number of the single bit set. In all other cases, that is, $2^k < x < 2^{k+1}$, the bits additionally set represent the binary fraction of the argument. This can be seen by writing

$$\begin{aligned} x &= \sum_{i=0}^{k} a_i 2^i = 2^k + \sum_{i=1}^{k} a_{k-i} 2^{k-i} \\ &= 2^k + 2^k \cdot \sum_{i=1}^{k} a_{k-i} 2^{-i} \\ &= 2^k \cdot \left[1 + \sum_{i=1}^{k} a_{k-i} 2^{-i}\right] = 2^k \cdot (1+z). \end{aligned} \tag{2}$$

Here again, k is the number of the uppermost bit set. The binary fraction z is always in the range $0 \leq z < 1$.

Taking the logarithm of the last expression yields $ld(x) = k + ld(1 + z)$, so the $ld(1 + z)$ also lies within the range $0 \leq ld(1 + z) < 1$. One method to compute the logarithm therefore requires

- splitting the argument into the uppermost bit set and the remainder,
- determining the number of the uppermost bit set,
- interpreting the remainder as a binary fraction with bit numbering relative to the uppermost bit,
- computing an approximate logarithm $ald(z) \cong ld(1 + z)$, which maps the range $0 < z < 1$ onto itself, and
- adding up the number of the uppermost bit set and the result of this mapping.

It should be noted that for an argument range of, say 64 bits, a result range of only 6 bits would be required for true integer result values. If computed in this way, the relative result error would be between $1/64 \cong 2$ percent and $1/1 = 100$ percent. A much higher precision is achieved if the result is scaled before truncation. This can be done by multiplying the result by a constant to exploit the full available integer width. For 64 bits the constant would be $2^{64 - ld(64)} = 2^{58}$. This multiplication can be done by shifting the result left an appropriate number of bits.

Mitchell[2] gave the simplest method for computing an approximate value of $ld(1 + z)$. He proposed using the approximation

$$\begin{aligned} ald(x) &= ald\left[2^k \cdot (1+z)\right] \\ &= k + ald(1+z) = k + z + f(z). \end{aligned} \tag{3}$$

This approximation corresponds to a linear interpolation between $ld(2^k) = k$ and $ld(2^{k+1}) = k + 1$. The computation requires only a normalization of the argument. If the uppermost bit set is shifted to the carry, the argument bit width minus the shift count equals k, and the remainder can just be added to the scaled value of k. This can be seen in the example shown in Figure 1 (a small bit width is taken for convenience).

The maximum error in this approximation is found at $z = 0.443$ and is equal to 0.087, that is, roughly 10 percent. For many applications this error is still acceptable. There, the simplicity of the described algorithm allows a very fast implementation in software. The routine displayed in Figure 2 computes $ald(x)$. The routine is written for a 68000 microprocessor (both personal computers mentioned above use this CPU) and assumes 32-bit, nonnegative integers for arguments. Results also have this range.

When this routine is executed on a 16-MHz CPU, it requires an execution time of maximally 43.6 μs (argument = 1) and minimally 6.1 μs (argument $\geq 2^{30}$). If the arguments are distributed randomly over the argument range, the routine requires an average computing time of 7 μs.

Due to the simplicity of Mitchell's approximation, it can easily be implemented in hardware. Jankowski-Tebe[3] described a logarithmic counter operating according to this method.

The approximation error can be reduced in a straightforward manner in two ways. Because with

Mitchell's method $ald(x)$ is always smaller than $ld(x)$, the maximum error can be reduced roughly by a factor of 2. This is done by adding a positive constant chosen to minimize $| ld(1 + z) - z |$ in the interval $z \in [0,1]$. Additionally, it is possible to make a linear approximation between all argument values 2^i as supporting points. But one possible improvement would be to use more supporting points (between the values 2^i) and to make a linear approximation between them. Combet et al.[4] and Hall et al.[5] suggested using the approximation $ald(1 + z) = az + bf(z) + c$ with proper values for a,b,c and a proper function f. Both methods reduce the maximum error of 0.086 (Mitchell); Combet's method reduces to a value of 0.013 and Hall's method, to a value of 0.017. We pay for the reduction with higher computational and/or hardware effort(s). Besides requiring at least four additions, a multiplication is also required, which takes about 30 times longer than a register addition (for a 68000 CPU).

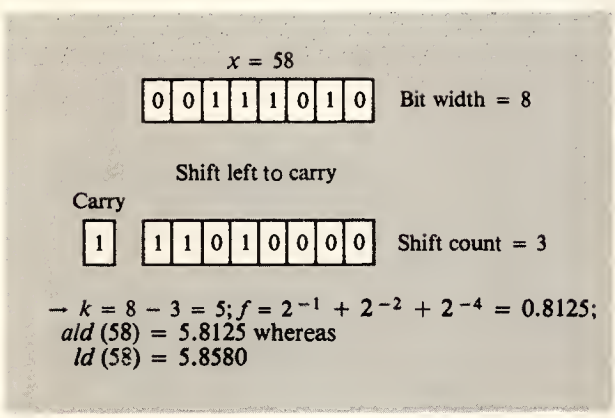


Figure 1. Principle of Mitchell's approximation.

Nonlinear approximation. The approximation error can clearly be reduced further if we use a nonlinear approximation. Marino[6] proposed a quadratic approximation that can be implemented in hardware. This method yields a maximum error of 0.004, which is smaller than Mitchell's error by a factor of 20. However, the usage of special-purpose hardware for the implementation of the logarithm is often not feasible, for example, for the personal computers mentioned earlier. A realization by software requires considerably more computational effort, even when compared to Combet's and Hall's approximations.

Lookup table with differential group PLAs. There are trade-offs between computational speed and approximation errors on one hand and between implementation in hard/software and the cost/performance ratio on the other hand. Lo and Aoki[7] proposed implementing the logarithm by special-purpose hardware using a programmable logic array (PLA) plus standard electronic devices such as adders. They first use the approximation

$$\begin{aligned} ald\,(x) &= ald\,[2^k \cdot (1+z)\,] \\ &= k + ald(1+z) = k + z + f\,(z), \end{aligned} \tag{4}$$

where $f(z)$ is an approximation to $ld(1 + z) - z$. This method was chosen, because $f(z)$ often has identical values for certain argument ranges. The idea is to group all argument values belonging to the same value of $ld(1 + z) - z$ and to assign each group to one product term of the PLA. However, the number of groups required is still comparable to the number of possible arguments. For a binary fraction with a width of 8 bits, 107 PLA groups are obtained instead of 256 groups using ROMs or RAMs. The maximum

```
* This procedure computes x = log2 (x) * 2**26
* using Mitchell's method
*
* Entry conditions:        D0.L = argument (>0)
* Return conditions:       D0.L = result
* D1 will be destroyed
*
LOG2:  MOVE.W    #31, D1     ; preset maximum loop counter
LOG2A: LSL.L     #1,D0       ; normalize argument by shifting
       DBCS      D1,LOG2A    ; uppermost bit to carry
       AND.B     #$C0,D0     ; clear lowest six bits
       OR.B      D1,D0       ; insert there number of upper-
                             ; most bit set
       ROR.L     #6,D0       ;rotate it to uppermost five bits
       RTS                   ; return
```

Figure 2. Routine for computing $ald(x)$ using Mitchell's approximation.

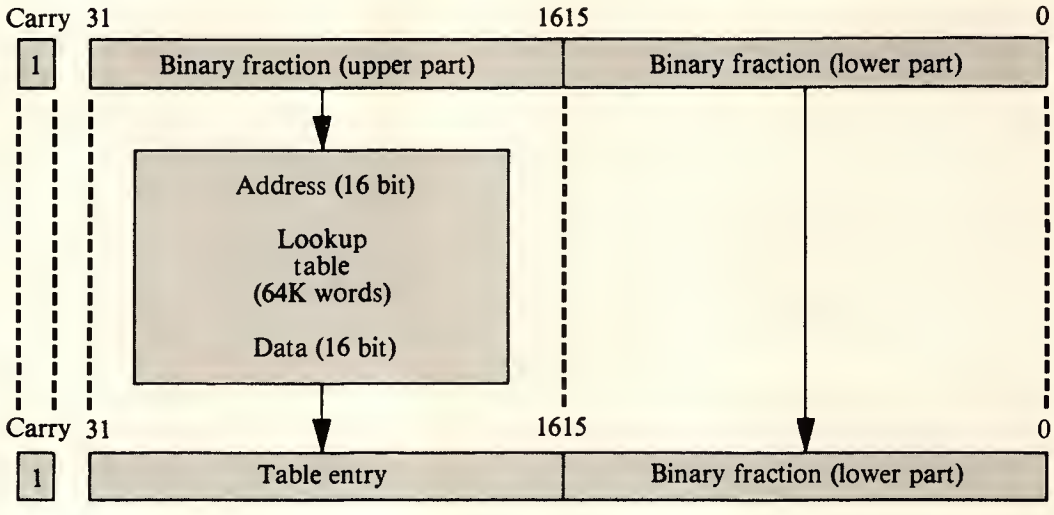


Figure 3. Principle of proposed approximation.

```
* This procedure computes x = log2(x) * 2**26
* using Mitchell's method combined with table lookup
*
* Entry conditions:      D0.L = argument (>0)
* Return conditions:     D0.L = result
* D1, D2 and A0 will be destroyed
LOG2:   LEA     Table, A0       ; setup pointer to lookup table
        MOVE.W  #31, D1         ; preset maximum loop counter
LOG2A:  LSL.L   #1, D0          ; normalize argument by shifting
        DBCS    D1,LOG2A        ;  uppermost bit to carry
        SWAP    D0              ; get high word of binary fraction
        CLR.L   D2              ; initialize high word to zero
        MOVE.W  D0,D2           ; copy high word of binary fraction
        ASL.L   #1,D2           ; make word offset for table access
        MOVE.W  0(A0,D2.L), D0; replace 1st 16 bit of binary
        SWAP    D0              ;  fraction with lookup table entry
        AND.B   #$C0,D0         ; clear lowest six bits
        OR.B    D1,D0           ; insert there number of upper-
                                ;  most bit set
        ROR.L   #6,D0           ; rotate it to uppermost five bits
        RTS                     ; return
```

Figure 4. Routine for computing $ald(x)$ using proposed approximation.

error in this example is 0.017, roughly the same as with the methods of Combet and Hall.

Combination method

I propose another combination of Mitchell's method and a lookup table. It can be implemented completely in software and is therefore well suited for usage in low-end personal computers. The basic idea is to split the binary fraction z into

$$z = \sum_{i=1}^{k} a_{k-i}2^{-i} = \sum_{i=1}^{m} a_{k-i}2^{-i} + \sum_{i=m+1}^{k} a_{k-i}2^{-i}, \quad (5)$$

where m is the number of bits used for addressing a lookup table. The approximation is then taken as

$$ald\ (1+z) = f(a_{k-1}, \ldots, a_{k-m}) + \sum_{i=m+1}^{k} a_{k-i}2^{-i}. \quad (6)$$

The function f is given as a lookup table. For all combinations of $a_{k-1}, \ldots, a_{k-m}$, the table contains one entry each with a value precomputed to minimize the approximation error. The operations required are

- normalizing the argument as done in Mitchell's method and
- replacing the first m bits of the binary fraction by an element of the lookup table.

Figure 3 shows operations for the case $m = 16$.

The value of m can be chosen arbitrarily. In principle, m might be anything from zero (Mitchell's method) to k (Brubaker's and Becker's method). This means that m can be optimized in regard to table size (2^m words, word width = m bits) and approximation error. In practice, one would choose m according to the available access width of the memory used, usually 8 or 16 bits.

It is easy to estimate the error made by this approximation. Without a lookup table, the result is the same as with Mitchell's algorithm. Using the first m bits of the binary fraction for a lookup table replaces these m bits by their exact, precomputed values. The only possible bit errors are in the bit positions for $m+1$ down to zero. If these bits were set to zero, the error would be $2^{-(m+1)}$. However, these bits are set according to the linear approximation $ald(1 + z) = z$. They therefore interpolate linearly between the exact values taken from the table. The error for these bits is therefore exactly the same as Mitchell's error scaled down by a factor of 2^{-m}. In the example given earlier, we have $n = 32$ and $m = 16$. The total error therefore equals $0.087 * 2^{-16} = 1.3 * 10^{-6}$. The routine shown in Figure 4 computes the binary logarithm in the proposed way.

When this routine is executed on a 16-MHz CPU, it requires an execution time of maximally 46.6 μs (argument = 1) and minimally 9.1 μs (argument $\geq 2^{30}$). If the arguments are distributed randomly over the argument range, the routine requires an average computing time of 10 μs.

The proposed algorithm to compute the binary logarithm has advantages when compared with earlier methods. Whereas some methods use special-purpose hardware to speed up computation or to reduce storage requirements, the new method can be implemented easily in software only. This capability is important, if a logarithm function is required in systems that do not use floating-point hardware and if special-purpose hardware add-ons are not feasible as with low-end personal computers.

I presented a procedure for a 68000 microprocessor, which is used in popular systems like the Macintosh and the Atari 1040 ST. This procedure is very simple – it uses 14 machine instructions only – and is computed quickly. For argument values distributed randomly over a 32-bit range, the average computation time is on the order of 10 μs for a 16-MHz processor, comparable to the speed of a floating-point coprocessor. Moreover, the approximation error is very small, depending on the actual implementation. For a 64K-word lookup table, this error is on the order of magnitude of 10^{-6} for the given example, a factor of 100 to 1000 smaller than the errors obtained with earlier approximations.

References

1. T.A. Brubaker and J.C. Becker, "Multiplication Using Logarithms Implemented with Read-Only Memory," *IEEE Trans. Comp.*, 1975, pp. 761-766.
2. J.N. Mitchell, Jr., "Computer Multiplication and Division Using Binary Logarithm," *IEEE Trans. Electr. Comp.*, 1962, pp. 512-517.
3. K. Jankowski-Tebe, "Logarithmus dualis fur Digitalsignale," *Elektronik*, 1983, pp. 99-100.
4. M. Combet, H. van Zonneveld, and L. Verbeck, "Computation of the Base Two Logarithm of Binary Numbers," *IEEE Trans. Elect. Comp.*, 1965, pp. 863-867.
5. E.L. Hall, D.D. Lynch, and S.J. Dwyer III, "Generation of Products and Quotients Using Approximate Binary Logarithm for Digital Filtering Applications," *IEEE Trans. Comp.*, 1970, pp. 97-105.
6. D. Marino, "New Algorithms for the Approximate Evaluation in Hardware of Binary Logarithm and Elementary Functions," *IEEE Trans. Comp.*, 1972, pp. 1416-1421.
7. H.-Y. Lo and Y. Aoki, "Generation of a Precise Binary Logarithm with Difference Grouping Programmable Logic Array," *IEEE Trans. Comp.*, 1985, pp. 681-691.



Reinhard Maenner joined the Physical Institute of the University of Heidelberg, West Germany, in 1975, after earning a Diplom in physics from the University of Munich. At the Institute, he developed a data acquisition system for large-scale experiments in nuclear physics. As part of this work, he wrote a real-time, demand-paged, virtual-memory multitasking operating system for a PDP-11/45. For these efforts he received the doctoral degree in 1979.

The author of over 25 papers on nuclear physics, computer architecture, and image processing, Maenner currently heads the group at the Physical Institute developing the Polyp multiprocessor. He is also a cofounder of Heidelberg Instruments, a company producing computerized optical instruments. He is a member of the ACM.

Questions about this article can be directed to the author at Physical Institute, University of Heidelberg, Philosophenweg 12, D-6900 Heidelberg, West Germany.

Reader Interest Survey

Indicate your interest in this article by circling the appropriate number on the Reader Interest Card.

Low 153 Medium 154 High 155

FEATURE

Effective Implementation of a Parallel Language on a Multiprocessor

Thomas L. Sterling, Albert J. Musciano, Ellery Y. Chan, and Douglas A. Thomae
Harris Corporation

A new multiprocessor execution environment integrates semantics of parallel control with mechanisms for synchronization of concurrent tasks.

Present technology provides computer designers with the potential for assembling a powerful computer at low cost by using multiple microprocessors to execute a single program. Such a machine is a member of the class of multiple-instruction stream, multiple-data stream (MIMD)[1] parallel computers called multiprocessors. These machines are distinguished by a number of roughly equivalent, tightly coupled processors.

Researchers have implemented several experimental multiprocessors, such as the CM* machine.[2] Unfortunately, more than a decade of active experimentation with single-program execution on multiprocessor architectures has produced few effective methods for achieving substantial performance gains. As a result, multiprocessors incorporating more than four processing elements have had little impact on commercial and scientific computing. We need advances in parallel-programming models and execution techniques to mate the abundance of inexpensive hardware to parallel applications.

A multiprocessor-based, parallel-program execution environment called SPOC (Simultaneous Pascal on Concert) provides a solution to the resource-utilization problem for a broad range of applications using small- and medium-scale multiprocessors. SPOC is based on a *concurrent thread* model of parallel computation. Each thread is a sequence of instructions that can be independently scheduled for execution.

Simultaneous Pascal[3] is a high-level language that provides the constructs to explicitly delineate and organize the concurrent threads of an application program. The *rendezvous control mechanism* (RCM) is the runtime strategy employed by SPOC to synchronize the termination of multiple active threads as an effective solution to the "join problem."

The Concert[4] multiprocessor testbed developed at MIT and Harris Corporation is the physical parallel-execution medium for programs written in Simultaneous Pascal (see Figure 1).

The parallel semantics of Simultaneous Pascal and the RCM were developed together, and each reflects the power and limitations of the other. SPOC is the synthesis of these two methods with the parallel resources of the Concert multiprocessor. It provides a complete execution environment for user applications.

Multiprocessor characteristics

From a physical standpoint, we can view a multiprocessor in the following terms:

- Number of processors,
- Processor power in terms of speed and word width,
- Memory size, speed, and distribution, and
- Communication interconnection for processors and memories.

0272-1732/87/1000-0046$01.00 © 1987 IEEE

This view is limited, since it conveys only the physical structure of a machine. Although we can deduce the peak performance level of the multiprocessor, this view gives no indication of the logical machine residing above the hardware, or its actual behavior.

The abstract model of a multiprocessor presents the machine as a collection of primitive functions that manage, process, and coordinate concurrent activities. We can characterize such a model in the following terms:

- Types of parallelism that can be exploited within an algorithm,
- Synchronization among concurrent activities,
- Context initialization and switching,
- Scheduling strategies for concurrent activities, and
- Locality of distributed objects and relevant activities.

The characteristics of a machine determine the *granularity* of the parallelism that can be used effectively by that machine. Granularity is the smallest amount of work that can be scheduled for execution without incurring significant performance losses due to overhead. We define overhead as the work performed by a multiprocessor to manage the execution of parallel activities that would not be performed by a uniprocessor executing the same program. Examples of overhead are synchronization of tasks, task scheduling, and context switching. A meaningful measure of the minimum-task granularity that can be efficiently employed is the amount of overhead incurred in support of the task.

Scalability is the maximum number of processors that can be efficiently used by a multiprocessor. This limit is largely determined by the system's granularity and performance degradation due to contention for access to shared resources. Scalability determines the most powerful multiprocessor that could be implemented with a particular architectural approach. The scalability of the architecture, combined with the parallel nature of a particular application program, defines the upper bound of the speed with which the program can be executed.

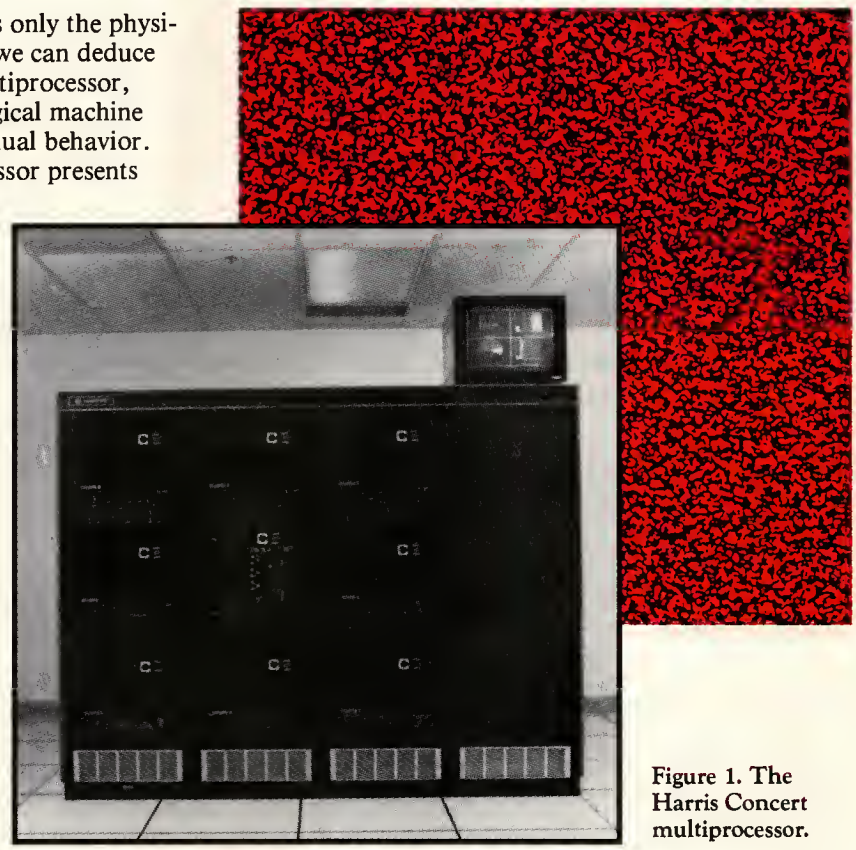


Figure 1. The Harris Concert multiprocessor.

Models of parallel synchronization

Several models of parallel computing exist, each imposing its unique synchronization requirements. Such models include multiprogramming, communicating sequential processes, and data flow.

Multiprogramming systems are very coarsely grained, running a number of programs at one time and using only the parallelism at the job boundaries. Although the multiprogramming approach exhibits limited scalability, Wulf and Bell employed it in the experimental C.mmp,[5] while Matelan used it in the commercial Flex/32.[6]

The Communicating Sequential Processes (CSP)[7] model is coarsely grained, with the available parallelism defined at the process level. Data passes between pairs of processes, either of which may be suspended until the other is ready to communicate. Here, too, scalability is limited, and the class of ap-

plications that may be readily captured with this model is restricted. To assist in this latter problem, researchers have developed several high-level languages that reflect this model, including Ada,[8] Occam,[9] and Concurrent Pascal.[10]

The data-flow computing model[11] uses fine-grained parallelism with data-driven synchronization. Its value-oriented paradigm limits the types of problems to which it is applicable; its fine-grainedness makes achieving an efficient multiprocessor implementation difficult.

Philosophy of parallel languages

Great debate surrounds the problem of writing parallel programs. Some parallel-programming languages require programmers to explicitly indicate the parallelism in their programs. Others require programmers to divide a program among several processors, and to provide synchronization and communication mechanisms. At the other end of the spectrum, smart compilers are used to derive the parallelism in a sequential program.

Many languages fall between these two extremes. A class of several languages reveals parallelism without programmer intervention. Unfortunately, these languages, such as data flow[12] and graph reduction,[13] require a purely value-oriented programming style and restrict the class of usable applications.

statement ::= ...
forall_statement
fork_statement
using_statement
lock_statement

forall_statement ::= **forall** *id* := *expr* **to** *expr* **do**
statement

fork_statement ::= **fork** *statement_list* **join**

using_statement ::= **using** *decl_list* **do**
statement

decl_list ::= *declaration* [; *decl_list*]

declaration ::= *id_list* : *type_id*

id_list ::= *id* [, *id_list*]

lock_statement ::= **locking** *expr* **do**
statement

Figure 2. Simultaneous Pascal parallel constructs.

The SPOC philosophy is to adopt a suitable sequential language and augment it with explicit parallel constructs. This philosophy encourages programmers to be conscious of the parallelism in their programs and helps them to write naturally parallel algorithms. Although this approach requires some initial learning, programmers quickly adapt to the new language features.

The SPOC model does not demand that programmers become intimately familiar with the underlying parallel machine. The burden of effectively mapping a parallel program onto the machine is left to the compiler and runtime system. In addition to freeing the programmer from this arduous task, the SPOC model promotes the creation of portable parallel code.

Simultaneous Pascal

The Simultaneous Pascal programming language is a superset of Standard Pascal.[14] Simultaneous Pascal contains all language features found in Standard Pascal and augments these features with several parallel-control constructs. These constructs allow the programmer to specify the portions of the program that can execute in parallel, to provide for a finer grained control of variable scoping, and to allow controlled access to global objects.

Simultaneous Pascal, unlike other parallel versions of Pascal such as Concurrent Pascal, promotes a thread-based model of parallel control. We can view a program as a collection of sequential pieces of code, or *threads*. Any single thread, once scheduled to execute upon a processor, will run to completion without interaction from any other processor. When a thread has finished executing, subsequent threads are then made available for execution. Parallel constructs in Simultaneous Pascal allow a programmer to indicate the threads in a program and to delineate the precedence constraints between various threads. The syntax of these new constructs is shown in Figure 2.

Three principal types of parallelism are available to the programmer in Simultaneous Pascal. Homogeneous parallelism allows the programmer to specify a single statement, many copies of which are to be executed in parallel. In heterogeneous parallelism, the programmer specifies several different statements to be executed in parallel. Operator-level parallelism allows the operands (which may be expressions) of an operator to be evaluated in parallel.

The Forall statement indicates homogeneous parallelism. The Forall statement allows the programmer to specify an index variable, a range of values over which the variable is to be instantiated, and a statement that is to be associated with each instance of the index variable. A simple Forall statement would be:

```
forall i := 1 to max do
  a[i] := b[i] + c[i];
```

In this example, *i* is the index variable, which will assume values over the range 1 to max. The statement

```
a[i] := b[i] + c[i]
```

will be executed max times in parallel, with each instance of the statement associated with a different value of *i*. Thus, parallel vector addition is being performed. All of the instances of the body of the Forall statement will finish executing before the statement after the Forall is scheduled for execution.

Heterogeneous parallelism is implemented with the Fork statement. The Fork statement allows the programmer to specify a list of statements, all of which may be executed in parallel. A simple Fork statement would be:

```
fork
  a := 1;
  b := 2;
  c := 3
join
```

The assignment statements may be executed concurrently. When all statements have completed, the statement following Join is scheduled for execution.

Operator-level parallelism allows the programmer to indicate those portions of an expression that can be evaluated in parallel. The operands of any binary operator can be executed in parallel by enclosing that operator within vertical bars. The following example uses operator-level parallelism:

```
r := sqr(sin(x))|*|sqr(cos(x));
```

Since calculating the square of the sine or cosine of a value is a relatively long process, the two subexpressions can be evaluated in parallel, reducing the time needed to evaluate the right-hand side of the Assignment statement. When both subexpressions have been evaluated, their product will be computed and assigned to *r*.

The lowest resolution of scoping in Standard Pascal is the procedure or function. That is, the *scope* of a procedure or function is the smallest portion of a program in which a name can be known. When writing parallel programs, the need often arises to have variables whose scope is local to the current thread. To meet this need, Simultaneous Pascal provides a new construct, Using, that allows finer grained scope control down to the level of a single statement.

The Using construct allows the programmer to specify a list of variables and a statement. The variables will only be known within the statement, and storage for those variables will only be allocated while the statement is executing. For example, the following code reverses the order of the elements of an array:

```
forall i := 1 to max div 2 do
  using temp : integer do
    begin
    temp := a[i];
    a[i] := a[max - i + 1];
    a[max - i + 1] := temp
    end
```

Each instance of the body of the Forall statement will be associated with a variable named *temp*. The storage allocated for temp will be local to each thread. Thus, each thread can assign to and use temp without interfering with other threads as they use their private temp variables. Without scope resolution at the statement level, note that the programmer would be required to create a temporary array so that each thread would have a unique location to use for intermediate storage. The naive approach, which is to simply declare temp in the scope containing the Forall statement, is incorrect; all threads will assign to and attempt to use the single instance of temp simultaneously, yielding incorrect results.

Simultaneous Pascal promotes a thread-based programming model. Unlike the CSP model of parallel processing, no way exists for a thread to send messages to other threads. However, the problem of having exclusive use of a global-data object is not eliminated. While CSP models would place the global object in some sort of monitor, and use semaphores to control entry into the monitor, Simultaneous Pascal does not have the ability to suspend threads waiting for access to the monitor.

Simultaneous Pascal does give the programmer the ability to serialize otherwise-parallel access to some global object by means of the Locking statement. The Locking statement allows the programmer to specify a lock variable and a statement to be executed if the lock is unlocked. The following code fragment shows how a thread would access the registers of some hardware device, guaranteeing that it has exclusive use of the device:

```
locking device.lock do
  begin
  device.control := write_data;
  device.data := 0;
  end
```

The variable named *device.lock* must be of type lock, which is unique to Simultaneous Pascal. Lock may be used like any other simple type, except that variables of type lock can only be passed as var parameters or used as the object of a Locking statement. Variables of type lock have only two values: locked and unlocked. Initially, a lock is set to the unlocked state. When the Locking statement is encountered, the thread will loop until the specified lock is in the unlocked state. The lock will then be placed into the locked state, and the thread will con-

tinue with the execution of the body of the Locking statement. When the body of the Locking statement has finished, the lock is returned to the unlocked state.

Note that threads *do not* suspend themselves when a locked lock is encountered. The processor on which the thread is executing will poll, waiting for the lock to become unlocked. Misusing locks will impair parallel-processor performance.

Runtime support mechanisms

SPOC presents many interesting implementation problems. The two most important involve (a) synchronizing several concurrent threads and (b) managing multiprocessor memory.

Rendezvous control mechanisms. A thread can be scheduled for execution when a specified set of preceding threads has run to completion. The difficulty of determining when all preceding threads have terminated is known as the *join problem*. Various computing models approach the join problem differently. These approaches can be distinguished by the amount of knowledge the system has concerning the location at which each join is to be performed. This *join location* (or *join point*) is a data structure in memory where the synchronization of multiple tasks is controlled.

Solutions to the join problem. At one extreme, special-purpose, synchronous systems such as systolic arrays require no join point, since hardware timing guarantees that join conditions will be satisfied when the operation is performed. At the other extreme, dynamic data-flow models require associative search techniques to match tokens destined for the same operation. The join location is not known until one of the tasks (in this case a single operation) has completed. Since the other task required for synchronization cannot know this location, a search must be made to locate the appropriate join point.

Some variants of the CSP model and the static data-flow model specify the join location at compile time. In the case of CSP, semaphores are defined to control the handshaking between concurrent processes. For static data flow, activity templates are defined at compile time that buffer and synchronize arriving values to determine when that template's operation is ready to be executed. There is no runtime cost in determining the join location, but the static program flow necessitated by the compile-time join locations can restrict the usefulness of these models.

SPOC uses the RCM to solve the join problem. (Our use of "rendezvous" bears no relationship to the concept of rendezvous in the Ada programming language.) The RCM allows SPOC to remember the join point of a thread when that thread is created. This method provides flexibility in programming, while eliminating the need to search for join points at runtime. We refer to this kind of join point determination as *runtime preassignment*.

In the Fork statement, the threads representing the various statements must join at the join keyword before program flow can exit the Fork statement. The threads created by the Forall statement must join before control passes on to the next statement. Thus, the parallel semantics of Simultaneous Pascal permit runtime preassignment of join locations for synchronization, preserving flexibility of application while minimizing synchronization overhead.

Rendezvous counters. Simultaneous Pascal synchronization requires that all threads within a particular parallel block complete before execution continues beyond the block. The order of completion is unimportant; so is the knowledge of which threads have not yet completed.

In a Fork statement, the number of threads to be joined is known at compile time. The number of threads to be joined in a Forall statement can be determined at runtime by evaluating the upper and lower bounds of the Forall range. In both cases, the number of joining threads is known before any of the threads is created.

Whenever threads are created, a small piece of memory is allocated to track the joining of those threads. This piece of memory is the join location for the threads. The join location contains a counter which is initialized to the number of threads to be joined. The address of the join location is passed to each thread as it is created. As a thread completes execution, it atomically decrements and reads the counter. If the value is greater than zero, the join is not yet complete. However, when the counter reaches zero, the join is complete, and program flow proceeds out of the current block.

This mechanism for synchronizing the completion of a set of threads is inexpensive, even when performed in software. It provides the functionality needed to resolve the join problem by exploiting the semantics of Simultaneous Pascal.

Dynamic structures of the RCM. One remaining issue is how the join address is transmitted to the threads it will synchronize. Consider the following simple example:

```
forall i := 1 to 10 do
  a[i] := b[i] * c[i];
```

Each instance of the Forall body is a single thread, and the address of the join location is passed to each thread when the thread is created. A more complicated example may pose a problem:

```
forall i := 1 to 10 do
  forall j := 1 to 10 do
    a[i,j] := b[i,j] * c[i,j];
```

In this case, the address of the join location representing the outer Forall is passed to each instance of its body. However, these threads immediately create new threads, to which is passed the address of the join location representing the inner Forall. Upon completion of the various instances of the Assignment statement, how will the join for the outermost join point occur?

SPOC employs a tree structure that grows and shrinks in response to the nesting of parallel blocks. Every join location also contains a pointer to the join location of the thread that created it. These pointers represent a tree, with each join location pointing to its parent. Each thread in the system has a pointer to some join point which is a leaf node of the tree. As parallel blocks are entered, the tree grows with the addition of new join locations. As these blocks join, the locations are removed and the tree shrinks.

When a join occurs, a thread is created to continue execution beyond the parallel block. This thread is given a pointer to the parent of the join point just synchronized. This new thread will, in turn, either extend the tree by entering a new parallel block or reduce the tree by participating in another join operation. Figure 3a shows the join location tree prior to the execution of a parallel thread. Figure 3b shows the tree with the additional join location for the executing threads. When they complete execution, the join location will be deallocated, and the tree will again look like Figure 3a.

The tree of join locations reflects the dynamics of program execution. Each thread need only retain a pointer to one join point in the tree, without regard to the depth of parallel nesting. The pointer to the relevant join location is passed to a thread when it is created. Relevant join information is stored within the structure; the expansion and contraction of the tree reveals this information at the correct times.

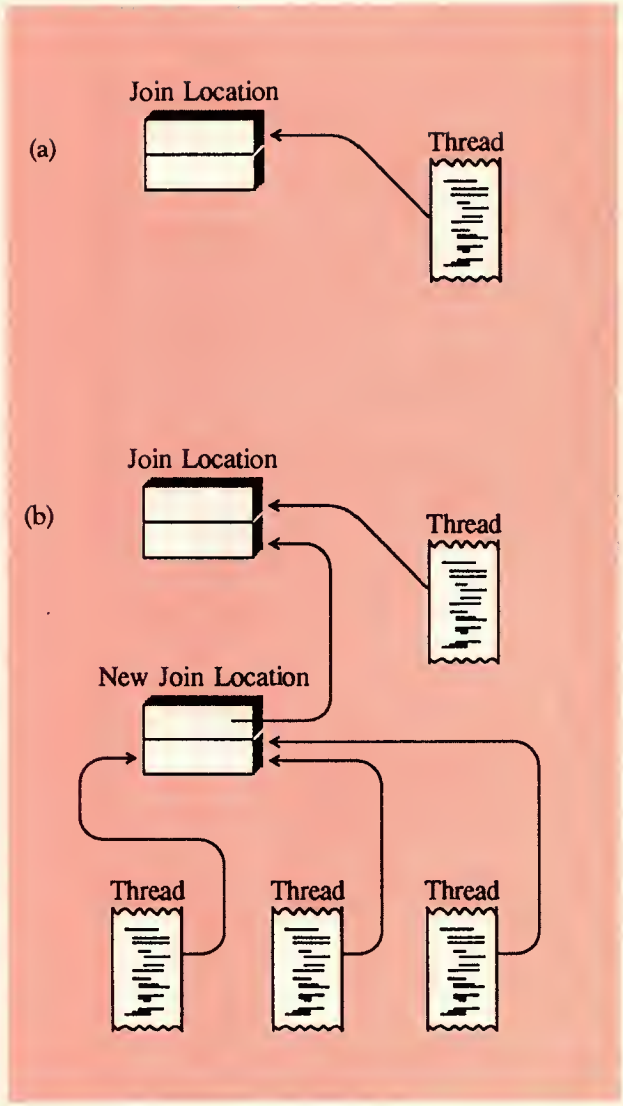


Figure 3. Join location before (a) and after (b) the Forall statement.

Memory management. The runtime needs of a Simultaneous Pascal program present a unique set of memory management problems for the runtime system. A program will require the standard Pascal heap, which is manipulated by the programmer using the New and Dispose intrinsic routines. In addition, a local stack mechanism is required so that runtime system routines can be accessed with minimal bus contention. Finally, each thread, procedure, or function requires a small piece of memory in which to retain any information local to that thread or routine. In order to meet these needs, memory is divided into three parts: the *heap*, the *local stacks*, and the *frame pools*.

Application program requests for dynamic memory allocation are satisfied from the heap. Any of several well-known algorithms for dynamic memory management can be used to implement the heap. All heap memory is located in the global memory space, so that pointers can be passed between processors.

Several libraries of routines in the runtime system handle I/O, arithmetic functions, system control, and the like. It is important that these routines run as fast as possible. Each processor in the system has a small stack located in the memory local to the processor. In order to reduce global-bus contention, parameters to these routines are pushed onto this stack whenever a library routine is called. In addition, these local stacks provide a place where processor interrupts can be fielded as quickly as possible.

Whenever a thread is created or a routine is called, a small piece of memory, called a *frame*, is associated with the thread or routine. For each thread, the frame contains back pointers to parent threads, the value of the thread index (if it was created by a Forall statement), and any local variables allocated within the thread by a Using statement. In the case of a routine, the frame contains both the static and dynamic back links for the routine, the return address parameters, and local variables. If a function is being called, space is also provided for a return value. Figure 4 shows the layout of the frames.

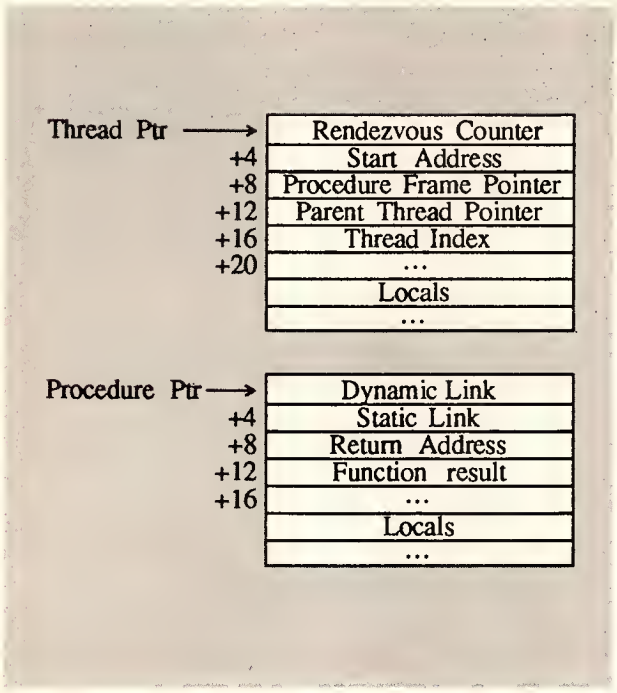


Figure 4. Layout of thread and procedure frames.

Several constraints are placed upon the frame pool. Many frames must be available, or the number of threads in the system at any moment will be reduced and the system will thrash for frames. The acquisition and release of a frame should occur quickly, since the speed of procedure calls and thread creation will be directly related to how quickly a frame can be made available.

The compiler and the runtime system jointly solve these problems. During compilation, the compiler determines the smallest usable frame size and passes this information to the runtime system. When the runtime system is initialized, a large chunk of memory is broken up into as many frames as possible, based on the size determined by the compiler. When a frame is required, it is removed from a list of free frames. When a frame is returned, it is placed back onto the free list. In order to reduce contention, the runtime system may allocate several disjoint frame pools, so that multiple processors can obtain frames simultaneously without contending for the same frame free list.

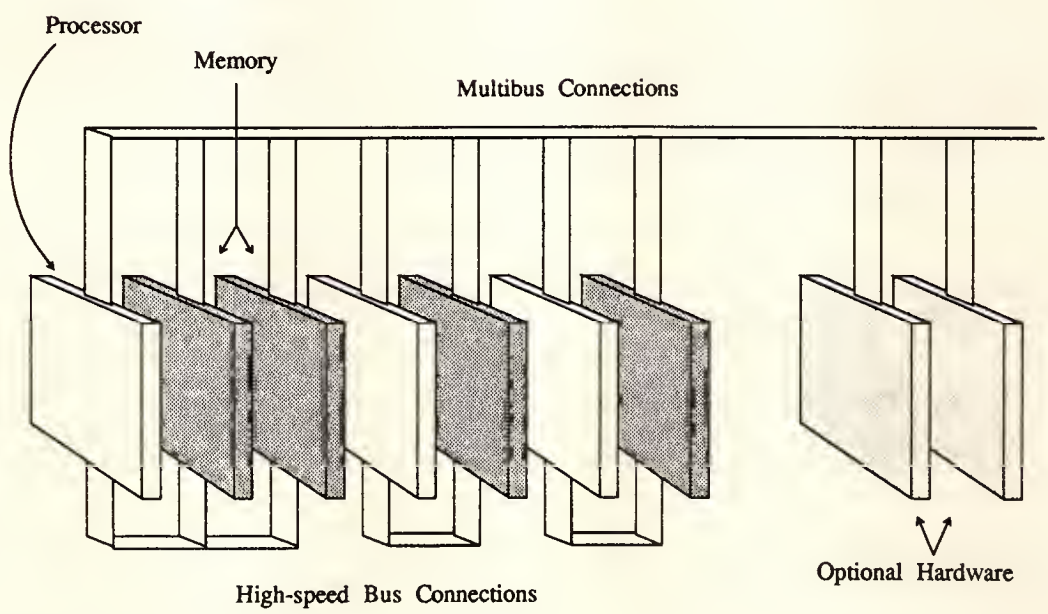


Figure 5. A Concert cluster.

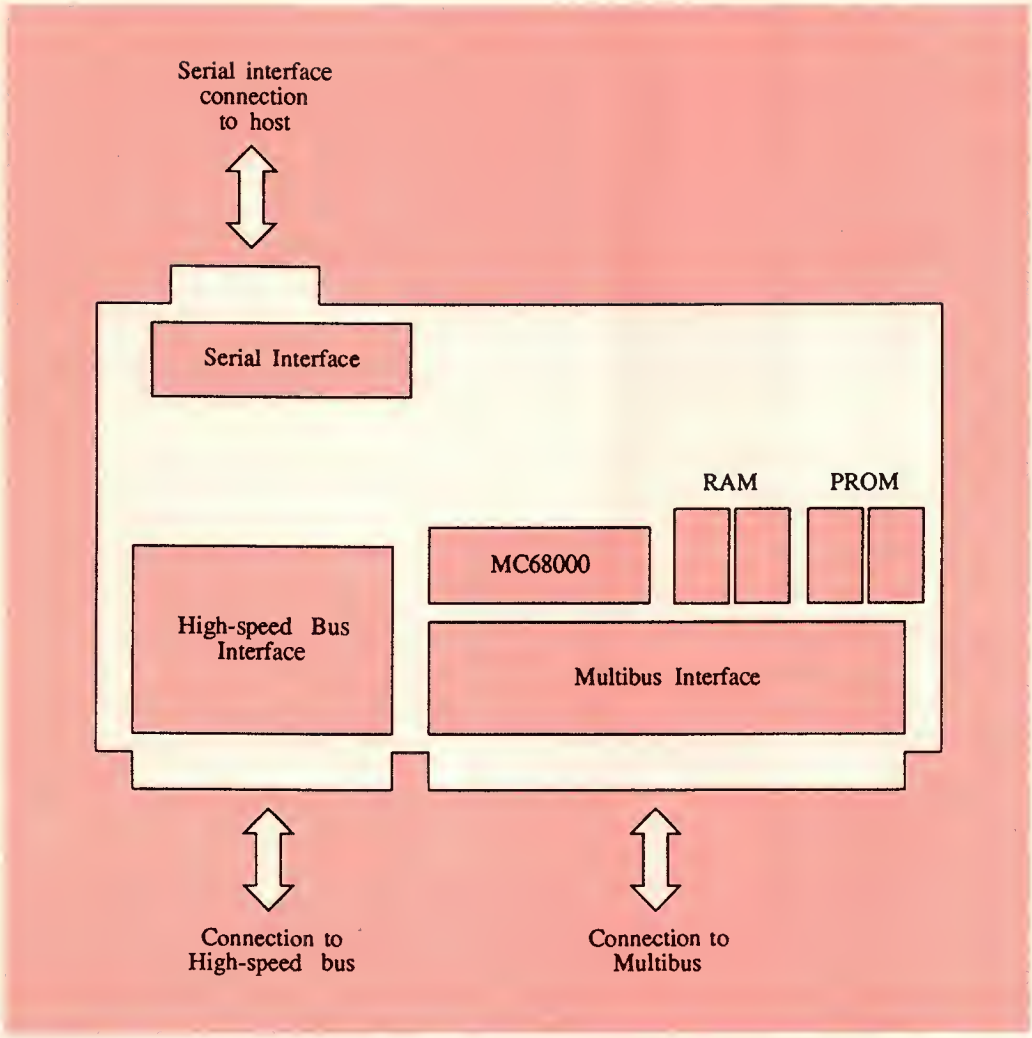


Figure 6. A processor module.

The Concert system

The Harris Concert multiprocessor is a tightly coupled, general-purpose computing resource that can support up to 64 processors. It differs from MIT's Concert machine[15] by using a crossbar switch rather than a ring bus interconnection scheme. Shared memory is intended to be the primary means of communication between processors. The Concert system hardware is shown in Figure 1.

The Concert structure exploits locality: processors and memory are grouped to decrease the cost of local accesses while (perhaps) increasing the cost of accesses to nonlocal resources. Groups of up to eight processors share a common backplane. Each backplane also supports several dual-ported memory modules, a global memory interface, and optionally a hard-disk controller, Ethernet controller, and/or custom system-monitoring hardware. This assembly is called a cluster (see Figure 5).

The cluster employs three levels of locality. Within a processor module (see Figure 6) the Motorola 68000 microprocessor has private access to 8K bytes of static RAM and/or PROM. The processor also has access via a private, high-speed bus to dual-ported RAM modules which share that bus. High-speed bus latency can be affected by Multibus accesses via the RAM module's other port, or by dynamic RAM refresh operations. The processor has access, using the Multibus, to the dual-ported RAM modules local to other processors in the cluster. The cost of this access can be increased by contention for the Multibus, by contention for the RAM module with its local processor, and by refresh cycles. Table 1 demonstrates various access costs and possible sources of contention.

A fourth level of memory hierarchy allows intercluster communication by providing a path from each cluster's Multibus, through a crossbar switch, to 8M bytes of interleaved global RAM (see Figure 7). The cost of accessing the global RAM can be affected by

Table 1.
Memory-access differences.

		Sources of contention			
RAM	Latency (ns)	Refresh	In same cluster	Other processor Using Multibus	In other cluster
High-speed bus	625	■	■		
On-board	750				
Multibus	1250	■	■	■	
Global	1500	■		■	■

Multibus contention, access to the same interleaved memory module by multiple clusters, and global RAM refresh cycles.

Results

The SPOC project will span several phases. The first phase, intended to prove the functionality of SPOC, is complete. We can currently create, compile, and execute Simultaneous Pascal programs on clusters of up to eight processors. Although our principal goal in the first phase was to achieve a functioning parallel processor, we also realized some performance gains. Programs which execute on several processors do exhibit performance enhancement as the number of processors increase.

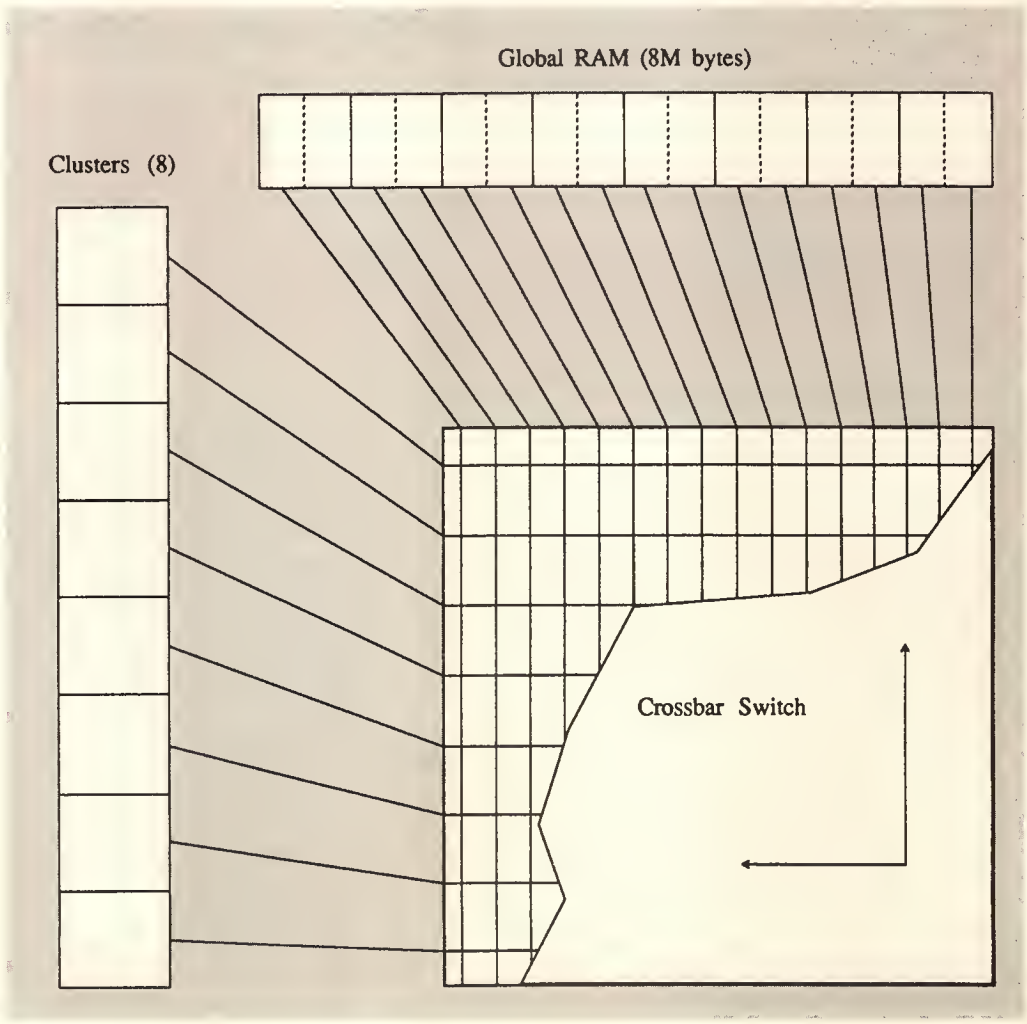


Figure 7. The Concert system.

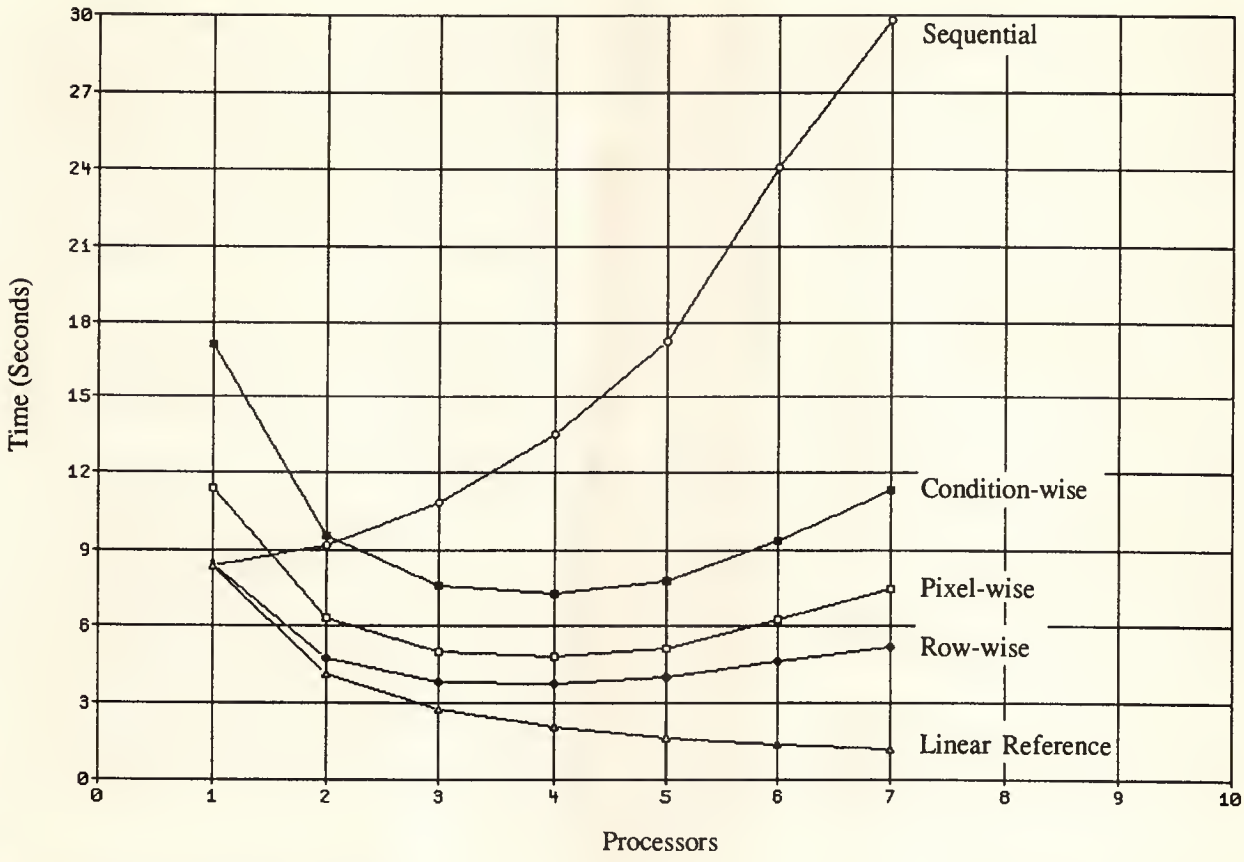


Figure 8. Digital-thinning performance.

We selected two applications to illustrate the effectiveness of the SPOC environment. The first is a digital-thinning algorithm,[16,17] used in optical applications to reduce complex images to their minimal, skeletal form. The second is Gaussian elimination, which is used in a variety of applications to solve linear equations.

Digital thinning. The digital-thinning algorithm is massively parallel and involves inspecting the neighbors of each pixel in an image. If the neighbors meet four different conditions, the pixel under consideration is allowed to remain in the image. Multiple passes are made over the image until it stabilizes and no more pixels can be removed.

The parallelism of this application is readily apparent. In all cases, each pixel of the image, and even each condition applied to each pixel, can be examined in parallel. Thus, the core of the algorithm can be described by the following fragment of code:

```
forall i := 1 to rows do
  forall j := 1 to columns do
    if image[i, j] then
      temp[i, j] :=
                neighbors(i, j)    |and|
                pattern(i, j)      |and|
                condition_c(i, j)  |and|
                condition_d(i, j);
image := temp;
```

Note the use of nested Forall statements to access each pixel of the image in parallel, as well as the parallel operator |and| to cause each of the conditions to be evaluated in parallel.

The performance graph of this algorithm in the SPOC system, relating execution time to the number of processors, is shown in Figure 8. The figure has several lines representing various thread granularities, a reference for the sequential version of the algorithm, and a reference representing the linear speedup expected of a perfect machine. A computing profile, which indicates the number of active threads at any given point, is shown in Figure 9. This profile indicates the parallelism available at each point during program execution.

The graph in Figure 8 gives tremendous insight into the performance of the SPOC system. Most interesting is the (rapidly!) climbing curve representing the sequential version of the program. When the sequential algorithm executes, one processor is performing

useful work while the remaining processors wait in vain for more threads to be created. These extra processors present a tremendous amount of global-bus contention, which hinders the progress of the sole working processor. This contention is particularly severe because the image being worked upon is held in global memory, along with the thread queue. Each atomic access to the thread queue by an idle processor precludes the use of the bus by the working processor. As more processors are added to the system, this contention increases.

Also of interest is the comparison between the versions of the program with different granularity. Because of its inherent parallelism, the digital-thinning algorithm is easily modified to be parallel in a row-, pixel-, or conditionwise manner. Currently, the amount of work performed for each condition, and thus for each pixel, is relatively small. Although each version of the algorithm shows some initial performance improvement when additional processors are added, the best scaling is achieved by the row-wise version, since it has the highest ratio of work performed versus scheduling cost.

Eventually, the addition of extra processors will begin to slow down the machine, as more time is spent waiting for and scheduling threads than is spent performing useful work. In addition, the extra bus contention created by extra processors, even though they are working, eventually bogs down the system.

Note that the image being thinned is only 52 × 48 pixels. Normal images in the optical realm would be orders of magnitude larger and would present significant amounts of work on a row-wise basis, providing for better scalability. Memory constraints within the current SPOC system prevented the use of a larger image.

Gaussian elimination. Gaussian elimination is a well-known technique for solving n linear equations in n unknowns. The algorithm consists of n - 1 phases, in which separate subphases pivot the matrix to bring the largest coefficient to the upper left, divide each element of the matrix by this coefficient, and finally subtract the upper row from each of the other rows. The pivoting phase is sequential in nature, while the division and subtraction phases can exploit either row- or element-wise parallelism. The code for each of these phases and subphases is shown in Figure 10. When all of these phases have been completed, a final series of parallel passes sets the upper triangle of the matrix to zero. The algorithm thus causes alternating periods of parallel and sequential activity.

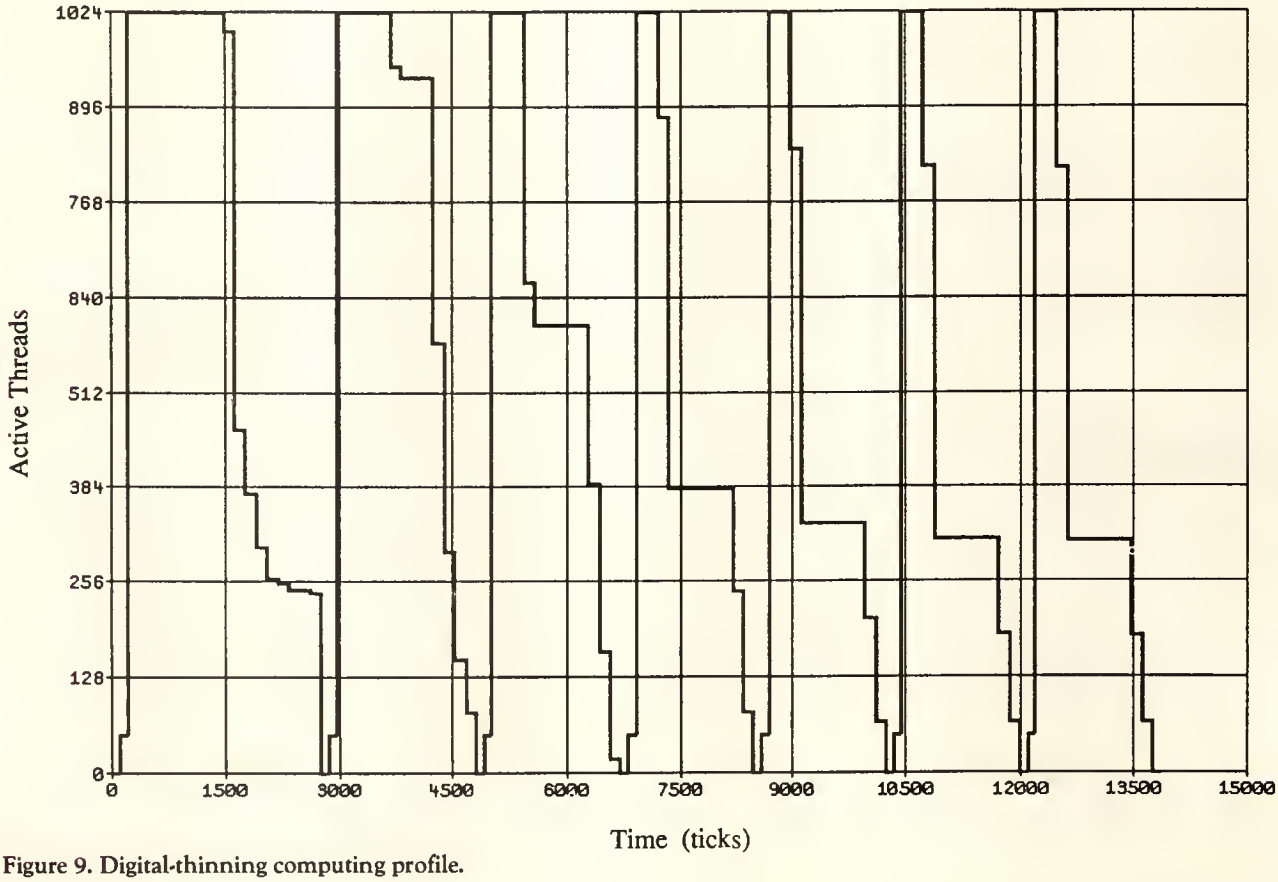


Figure 9. Digital-thinning computing profile.

```
for submat := 1 to size - 1 do
  begin
  { pivot the matrix }
  for i := submat + 1 to size do
    if abs(a[submat, column[i]]) > abs(a[submat, column[submat]]) then
      begin
      temp := column[submat];
      column[submat] := column[i];
      column[i] := temp
      end;
  { divide each row by the value at the upper left corner }
  forall i := submat to size do
    if (a[i, column[submat]]<>0.0) and (a[i, column[submat]]<>1.0) then
      forall j := submat + 1 to size + 1 do
        a[i, column[j]] := a[i, column[j]] / a[i, column[submat]];
  { subtract the top row from each of the others }
  forall i := submat + 1 to size do
    if a[i, column[submat]] <> 0.0 then
      forall j := submat to size + 1 do
        a[i, column[j]] := a[i, column[j]] - a[submat, column[j]];
  end
```

Figure 10. Simultaneous Pascal version of Gaussian elimination.

We tested two versions of the program, one exploiting element-wise parallelism, the other using row-wise parallelism. The performance graph and computing profile are shown in Figures 11 and 12.

Again, the sequential version suffers from bus contention in a manner similar to the digital-thinning algorithm. The losses are not so dramatic, however, because the Gaussian elimination program spends most of its time performing floating-point arithmetic. The library routines execute only in local memory and are not impacted by global-bus contention. Contention does occur, though, when a processor is fetching operands to pass in to the library routines.

The scalability of the Gaussian elimination program is much better than the digital-thinning program, because the smallest thread has more work to perform. Thus, the ratio of useful work to scheduling costs rises, and the scalability improves.

As in digital thinning, the row-wise version of the program scaled best and was still improving at seven processors. It most likely would have bottomed out around eight or nine processors.

Observations from experience. Although the previous examples show some weakness in the current implementation of SPOC, we obtained encouraging results even before we made performance modifications to the system. The fact that the system shows near-linear speedup for a small number of processors indicates that the goal of moderate scalability is obtainable.

Significant performance losses are occurring in at least two areas within SPOC. The first is global-bus contention caused by multiple processors attempting to gain exclusive access to the various thread and frame queues in the system. The performance curves of the two example sequential versions in Figures 8 and 11 demonstrate the effects of the unoptimized, runtime system software. The second source of performance loss is the overhead cost of scheduling a thread. Different performance curves representing various levels of granularity illustrate the impact of this overhead in Figures 8 and 11. The coarsely grained versions of each application ran more quickly than their finely grained counterparts.

The primary goal of the next phase of SPOC is to enhance the system's performance characteristics. This includes finding new mechanisms to prevent the bus and queue contention now occurring in the system, as well as reducing the time required to create and schedule a thread. Although our current efforts are directed toward software solutions to these problems, eventually hardware support for these operations may provide the best answer.

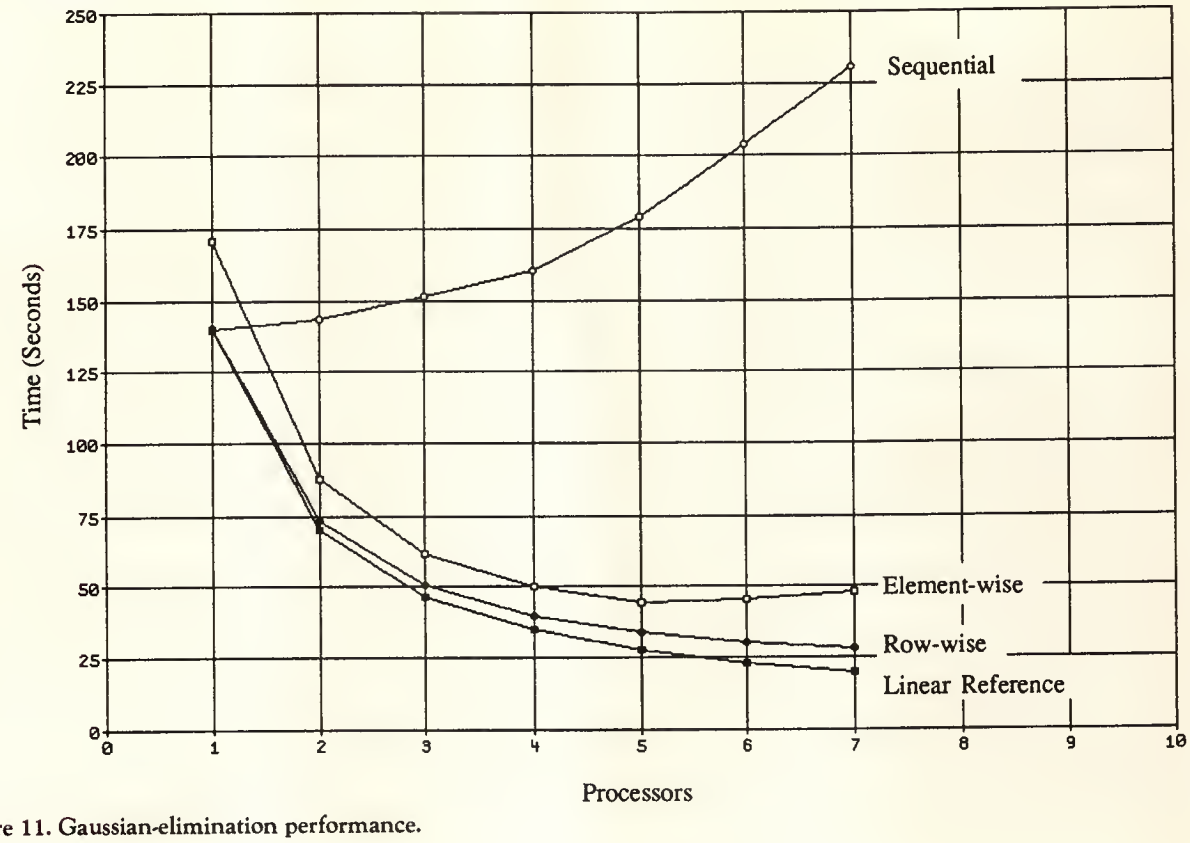


Figure 11. Gaussian-elimination performance.

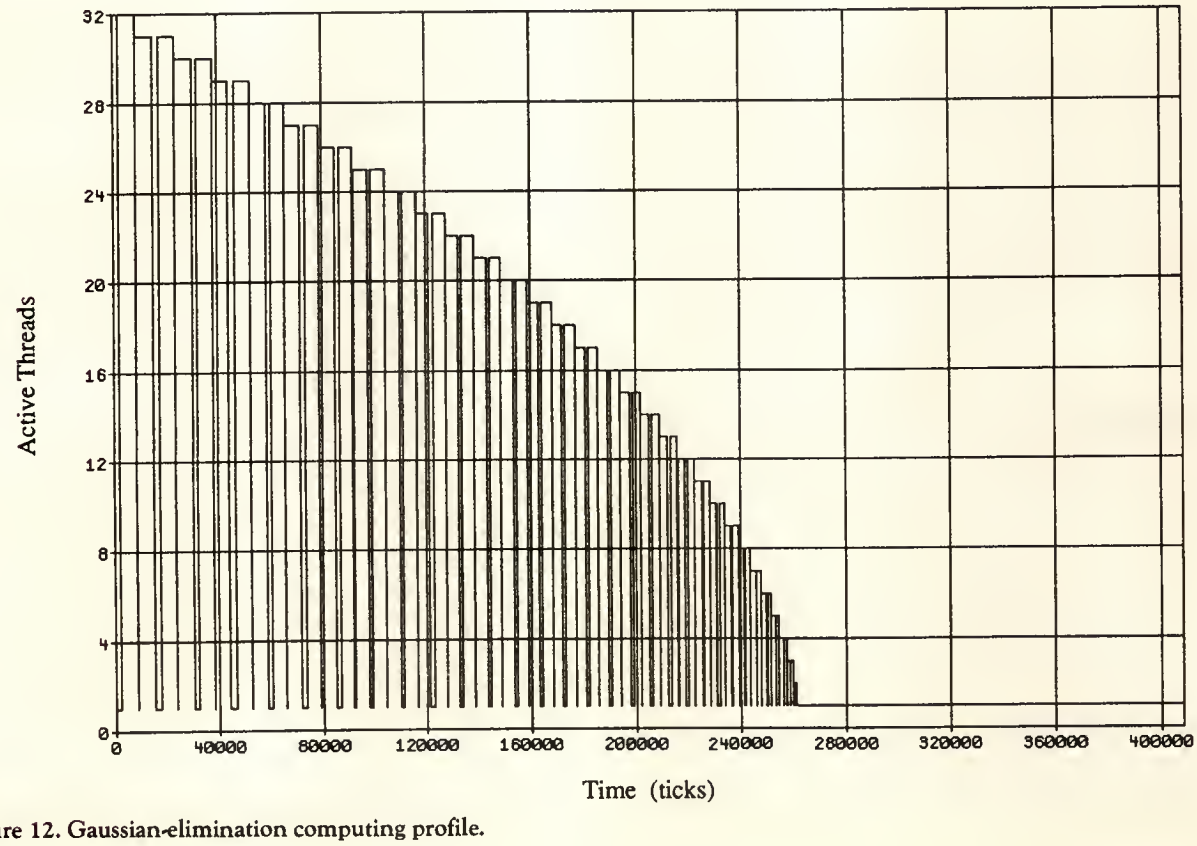


Figure 12. Gaussian-elimination computing profile.

The SPOC user environment

The ultimate value of a multiprocessor is dependent not only on how fast it executes a program but also on how easily the application programs can be created, debugged, and run. The SPOC multiprocessor environment includes an array of tools to facilitate rapid program development and support convenient user interaction. This environment encompasses both the Concert multiprocessor and its DEC VAX 11/750 host with Unix Berkeley 4.2, providing a rich set of tools for program development. It includes the host's Simultaneous Pascal software development environment, Concert's Simultaneous Pascal program execution environment, and the SPOC user interface.

The host Simultaneous Pascal environment. The host machine provides several tools to assist in the development of parallel programs. These tools include a compiler to generate serial versions of parallel programs suitable for host emulation, along with a graphics generator to produce computing profiles for the emulated code. Other programs are available in the host environment to detect parallel semantic errors, such as race conditions when accessing shared variables. Finally, the host machine contains the Simultaneous Pascal cross-compiler to generate parallel code suitable for execution on the Concert multiprocessor.

The Concert Simultaneous Pascal environment. The Simultaneous Pascal program execution environment on the Concert multiprocessor consists of four major subsystems: the SPOC user interface, the Concert runtime system, the debugger, and hardware instrumentation for performance analysis.

The SPOC user interface. The SPOC user interface (Figure 13) provides users with their primary impressions of the system. It performs three functions:

- It removes the burden of system housekeeping and initialization from the experimenter, who may be unfamiliar with the system.
- It performs file transfers and session recording at the user's request.
- It supports debugging tools which are invoked at the user's request or when a fault occurs on a processor in the SPOC configuration.

The available debugging tools include many low-level functions needed by the SPOC system implementers, who use the same user interface while doing SPOC development work. As the system matures, we expect the implementers to become experimenters themselves. The unified environment provided by the tools will allow the experimenters to move easily between implementing new features and evaluating changes in performance. The implementers' experience and the immediate feedback to their own work may well result in some of the most direct and valuable parallel-processing insight SPOC will produce.

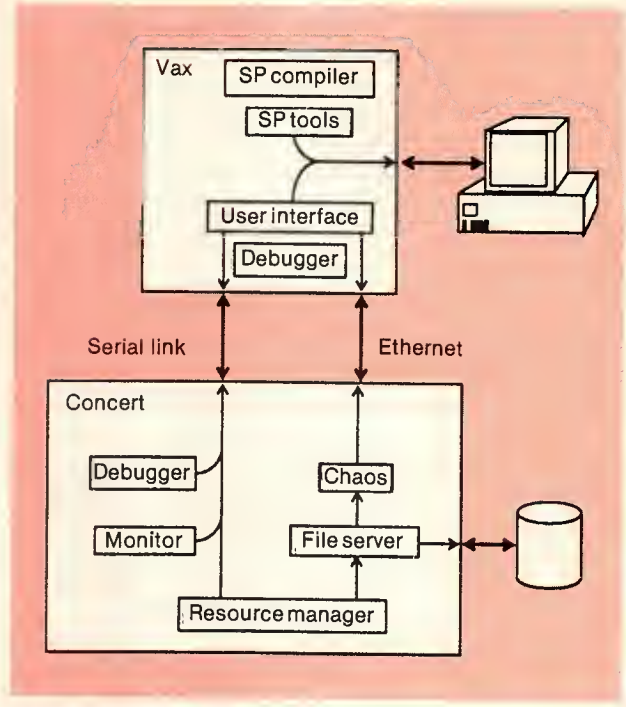


Figure 13. The SPOC user interface.

Limits of Simultaneous Pascal semantics

Simultaneous Pascal is an effective tool for representing many algorithms in a way that captures much of their available parallelism. The language and the underlying RCM are proving to be practical except for the following deficiencies.

List operations. The Forall construct permits the same action to be performed in parallel on all elements of arrays and sets. Currently, however, a Forall-like operation cannot be performed on the elements of a linked list. Because list processing is important for symbolic computation, the inability to apply a parallel operator to such data structures is an unfortunate limitation. Difficulty arises from the dynamic character of lists: Their length can only be determined by tracing pointers.

It would be desirable to augment Simultaneous Pascal with a construct that allowed parallel operations over lists and other dynamic data structures. Such a construct would allow the programmer to traverse a data structure, creating a thread for each element of the structure. A future version of Simultaneous Pascal may include this feature.

Overly constrained structure. Simultaneous Pascal encourages a programming style that promotes *chunky* parallelism resulting in hourglass-like computing profiles. Precedence constraints for a thread are often satisfied when only a few of the threads in a preceding parallel statement have completed. Since all the threads in a parallel block must terminate before any of the succeeding threads may begin, the execution of subsequent threads is unnecessarily delayed. This behavior can restrict the parallelism available in an algorithm. The impact of these constraints on execution behavior merits further study.

Artificial data dependencies. Program flow in Simultaneous Pascal is tied directly to the control relationships expressed in the program. Although this is satisfactory for many applications, the control-flow model can be artificially constraining. The Future construct in MultiLisp[18] provides a means for removing unnecessary program-flow constraints, leaving only the critical data dependencies. If a primitive operation needs an operand value before it has been computed, the thread is suspended. Once the value has been determined, the old thread is reawakened.

This form of parallelism is appropriate in symbolic applications where large, compound-data structures are manipulated in an unpredictable manner. Data-manipulation primitives that do not inspect the value will not be blocked on an uncomputed value. Programming with partially computed values, as in building data structures in which only some of the values have been calculated, becomes possible. This data-driven synchronization provides many opportunities for overlapped computation that cannot be captured in Simultaneous Pascal.

Granularity of parallelism. While too little algorithmic parallelism results in poor efficiency due to low processor utilization, too much parallelism can be equally detrimental to multiprocessor performance. Having too few processors results in the serialization of potentially parallel threads, while retaining the overhead associated with those threads. This serialization can substantially reduce the effective performance of a medium-scale system. A technique called *aggregation* would allow the compiler to build fewer, more coarsely grained threads from many finely grained ones. Aggregation would eliminate excessive overhead losses while maintaining sufficient granularity to retain good load-balancing.



A second problem is a parallelism explosion which reveals so much parallelism at one time that the system's capacity to keep track of all the pending threads is saturated. Representing large groups of similar threads with a single object in the system can minimize this occurrence. As processors request threads to execute, small numbers of threads can be created, with their (as yet) uninstantiated siblings represented by the object.

Unfair thread assignment. An objective of Simultaneous Pascal and SPOC is to eliminate the overhead costs of context switching incurred by suspending tasks. This approach is in sharp contrast to some CSP-based multiprocessors, the use of I-structures in dynamic data flow, and the use of Future in MultiLisp. A characteristic of Simultaneous Pascal is to assume that a thread, once scheduled for execution on a specified processor, will reside on that processor until it has completed. Processors are not time-sliced among active threads. This characteristic constrains the style of programming and eliminates the use of mutual exclusion as a means of interthread synchronization. All precedence constraints of a thread must be satisfied prior to the thread being scheduled for execution. A negative consequence of this unfair scheduling policy is that deadlocks can occur if CSP-style synchronization of communication between two or more threads is attempted.

Current state of SPOC development. We have achieved our goal of assembling a system that will accept Simultaneous Pascal source code, compile it to run on a sequential or a parallel machine, cause it to be executed on the desired machine, and produce statistics indicating the effectiveness of the program. We plan to continue enhancing performance. The current system is intended as a functional proof of concept and is not a high-performance vehicle. As SPOC performance improves, the techniques that evolve will determine new approaches to the design of parallel system hardware.

This last result is of particular importance to the parallel processor architect. Multiprocessors are implemented with off-the-shelf microprocessors because they are readily available. These processors are not designed to be constituent elements of large multiprocessors, and as a result provide little support for parallel processing. SPOC can help determine the nature of that support and the potential advantage of having that capability available in hardware instead of in software. The next generation of multiprocessors will include special hardware to efficiently control parallel execution.



Efficient control of parallel activities remains one of the critical unresolved problems inhibiting the application of multiprocessors to general-purpose computing. The SPOC multiprocessor execution environment embodies an experimental approach that integrates semantics of parallel control with mechanisms for synchronization of concurrent tasks on a medium-scale multi-microprocessor.

SPOC executes programs written in the experimental parallel programming language, Simultaneous Pascal. The parallel control constructs in Simultaneous Pascal are the Forall and Fork statements. Also, parallel operators permit subexpressions to be evaluated concurrently. The Locking statement enables exclusive access to shared data objects. The Using statement allows fine-grained scoping of temporary variables without resorting to procedure calls.

SPOC synchronizes the termination of concurrent threads of a parallel statement with the RCM. When a parallel statement is executed, a counter is allocated and used to keep track of the number of terminating threads, providing a means for synchronization. Since parallel statements can be nested, all the counters are linked in a tree structure that reflects the dynamic state of statement nesting. The cost of managing the counters is low, even when implemented in software. The RCM fully supports the synchronization requirements dictated by the parallel-control semantics of Simultaneous Pascal.

SPOC programs execute on the Concert multiprocessor. It incorporates 64 16-bit, MC68000 microprocessors, each with 512K bytes of local memory, and 8M bytes of shared memory. The machine is divided into eight clusters, each containing eight processors. Each cluster has a high-speed common bus and is connected to the global memory by means of a crossbar switch. Additional global registers provide for system-wide interrupts. SPOC is integrated with a host VAX 11/750 via Ethernet. Cross-compiler, debugger, and user-interface tools reside on the host and provide easy access to SPOC facilities.

The SPOC project is a relatively new, yet fully functional, system. Current activities include optimizing the runtime system code to minimize losses due to shared resource contention and extending the base of application programs used to test SPOC. Preliminary results show acceptable performance gains for small numbers of processors. We anticipate that future enhancements to SPOC will improve scalability, allowing full use of Concert multiprocessor resources.

Acknowledgments

We wish to recognize the contributions made by our colleague Michael D. Noakes in the formulation of the initial SPOC concept and in the implementation the SPOC resource manager. We also wish to acknowledge the leadership and guidance provided by Robert H. Halstead of MIT during the initial development of Simultaneous Pascal and the Concert multiprocessor. Finally, we wish to convey our thanks to our colleagues in the Advanced Technology Department at Harris Corporation for their many valuable suggestions during the creation of SPOC.

References

1. M.J. Flynn, "Some Computer Organizations and Their Effectiveness," *IEEE Trans. Computers*, Sept. 1972, pp. 948-960.
2. R.J. Swan, S.H. Fuller, and D.P. Siewiorek, "Cm*—A Modular, Multi-Microprocessor," *AFIPS 46, Proc. 1977 NCC*, AFIPS Press, Montvale, N.J., pp. 637-644.
3. T.L. Sterling, "Parallel Control Flow Mechanisms for Dynamic Scheduling of Tightly Coupled Multiprocessors," PhD thesis, EECS Dept., MIT, Cambridge, Mass., May 1984, pp. 24-29.
4. T.L. Anderson, "The Design of a Multiprocessor Development System," Tech. Report MIT/LCS/TR-279, Lab. Computer Science, MIT, Sept. 1982.
5. W.A. Wulf and G.C. Bell, "C.mmp – A Multi-mini-processor," *Proc. AFIPS 1972 Fall Jt. Computer Conf.*, AFIPS Press, 1972, pp. 765-777.
6. N. Matelan, "The FLEX/32 Multicomputer," *12th Ann. Int'l Symp. Computer Architecture*, CS-IEEE, Los Alamitos, Calif., June 1985, pp. 209-213.
7. C.A.R. Hoare, "Communicating Sequential Processes," *Comm. ACM 21*, Prentice-Hall Int'l, London, England, Aug. 1978, pp. 666-677.
8. N. Gehani, *Ada Concurrent Programming*, Prentice-Hall, Inc., New York, N.Y., 1984.
9. Inmos Ltd., *Occam Programming Manual*, Prentice-Hall Int'l, 1984.
10. P. B. Hansen, *The Architecture of Concurrent Programs*, Prentice-Hall, Inc., 1977.
11. J.B. Dennis and D.P. Misunas, "A Preliminary Architecture for a Basic Data-Flow Processor," *Proc. Second Ann. Symp. Computer Architecture*, CS-IEEE, Los Alamitos, Calif., Dec. 1974, pp. 126-132.

12. W.B. Ackerman, "Data Flow Languages," *Computer*, Feb. 1982, pp. 15-36.

13. J. Backus, "Can Programming be Liberated from the von Neumann Style? A Functional Style and Its Algebra of Programs," *Comm. ACM*, ACM, New York, N.Y., Aug. 1978, pp. 613-641.

14. D. Cooper, *Standard Pascal User Reference Manual*, W.W. Norton & Company, New York, N.Y., 1983.

15. R.H. Halstead et al., "Concert: Design of a Multiprocessor Development System," *13th Ann. Symp. Computer Architectures*, CS-IEEE, Los Alamitos, Calif., June 1986, pp. 40-48.

16. T.Y. Zhang and C.Y. Suen, "A Fast Parallel Algorithm for Thinning Digital Patterns," *Comm. ACM 27*, ACM, Mar. 1984, pp. 236-239.

17. H.E. Lu and P.S.P. Wang, "A Comment on 'A Fast Parallel Algorithm for Thinning Digital Patterns,'" *Comm. ACM* 29, ACM, Mar. 1986, pp. 239-242.

18. R.H. Halstead, "Multilisp: A Language for Concurrent Symbolic Computation," *ACM Trans. Programming Languages and Systems*, ACM, Oct. 1985, pp. 501-538.



Thomas L. Sterling is a senior principal engineer in the Advanced Technology Department of Harris Corporation's Government Systems Sector. His interests include parallel-control flow models of computation, high-level language parallelism constructs, static data-flow architectures, and computational dynamics.

Sterling received his PhD in electrical engineering from the Massachusetts Institute of Technology in 1984.



Ellery Y. Chan is a lead engineer in the Advanced Technology Department of Harris Corporation's Government Systems Sector. Prior to 1984, he was a hardware designer at Atex Incorporated. His current research interests include computer graphics, parallel processing, and user-interface design.

Chan received a BS in computer science and engineering from the Massachusetts Institute of Technology in 1981.



Albert J. Musciano is a lead software engineer in the Advanced Technology Department of Harris Corporation's Government Systems Sector and has been with Harris since 1982. His current research interests include parallel processing, language and compiler development, user-interface design, and software tools development.

Musciano received a BS in information and computer science from the Georgia Institute of Technology in 1982.



Douglas A. Thomae is a lead engineer in the Advanced Technology Department of Harris Corporation's Government Systems Sector and has been with Harris since 1982. His current research interests include parallel architectures and functional languages.

Thomae received a BS in electrical engineering from the University of Wisconsin-Madison in 1982.

Questions about this article can be directed to Musciano at the Harris Corporation, PO Box 37, MS 3A/1912, Melbourne, FL 32902, or via electronic mail to chuck@harris-trantor.arpa.

Reader Interest Survey

Indicate your interest in this article by circling the appropriate number on the Reader Interest Card.

Low 159 Medium 160 High 161

An Application-Specific Coprocessor for High-Speed Cellular Logic Operations

Robert E. Jenkins and D. Gilbert Lee, Jr.
Johns Hopkins University Applied Physics Laboratory

Why develop a special coprocessor to perform one lengthy computation? Performance is much better than with equivalent software and fast-turnaround development times will soon be possible.

The maturing technology in full-custom VLSI chips significantly affects the alternative approaches available for processing tasks. Application-specific chips, designed and fabricated inexpensively and quickly, often provide a viable alternative to software implemented on a general-purpose mainframe or a microcomputer. The use of custom VLSI chips makes a complex processor small enough and its power consumption low enough to be comfortably hosted by a personal computer. The personal computer provides a nicely suited front end for interactive use of the processor. This combination is a natural extension of the old idea of an attached array processor for matrix computations or even a floating-point coprocessor chip for PC-sized systems. The difference is that with custom VLSI one can think in terms of a high-performance coprocessor designed for a quite narrowly defined computing problem.

Does it make sense to develop a special coprocessor to perform, say, a specific Kalman filter computation and nothing else, rather than write a software version for a general-purpose computer? It does, but only if (1) the chip can be produced in about the same order of magnitude of time and cost as needed to develop equivalent software, and (2) it has a marked improvement in performance over software.

Chips designed with silicon compilers and fabricated via a fast-turnaround foundry service are beginning to make the application-specific coprocessor a sensible idea for a computing task that does not have a wide commercial market. At Johns Hopkins we have been exploring this potential through example projects to develop some narrowly defined processors at the very lowest end of the cost spectrum and then assess the usefulness of the results. Ideal candidates for such an approach are well-defined computations that will be run a large number of times—and which are formidable enough to cause the computing time on a programmable machine to be troublesome.

One of our selected computational tasks for an application-specific coprocessor is the cellular logic operation, or CLO, which is a special case of cellular automaton processing. Cellular automata are dynamical systems (introduced in 1948 by J. von Neumann and S. Ulam) made up of a regular array of sites, or cells, which can each be in one of a finite number of states. The system evolves to new states in discrete time steps through the interaction of each cell with its nearest neighbors. Such sys-

0272-1732/87/1000-0063$01.00 © 1987 IEEE

tems have been studied in full generality. However, most practical applications thus far have dealt with a two-dimensional array or grid, with a small number of possible states for the cells. In most such practical systems the state of a cell at time $t+1$ depends only on its state and those of its eight neighbors at time t. t can be expressed as the function:

$$X_{i,j}^{t+1} = F\left(X_{i,j}^{t}, X_{i+1,j}^{t}, X_{i-1,j}^{t}, \ldots, X_{i,j-1}^{t}\right). \quad (1)$$

The notation

$$X_{i,j}^{t} \quad (2)$$

denotes the state of site (i,j) at time t.

If the number of possible states is two, we have the special case of a binary array, where the function F becomes a 9-bit Boolean mapping referred to as a cellular logic operation. One defines this mapping by completely specifying the table of 2^9 possible states in the function. This binary form of automata represents a black-white image and has been used in this form for extensive image processing applications.[1,2] This special case is also the form for the well-known Game of Life neighborhood rule, invented by J. Conway in 1970, in which the individual cells live and die as a result of the number of neighbors.[3]

Research in the application of cellular automata has been ongoing for some time and is more active than ever. Interesting and promising new approaches based on cellular automata models have been proposed in a wide variety of physical and biological modeling problems, ranging in nature from galactic structure[4] to DNA sequences.[5] A representative sample of the ongoing research may be found in Farmer et al.,[6] which includes a good overview by S. Wolfram in the preface. Applications to image analysis are very well developed.[1,2] Preston at Carnegie Mellon and his collaborators at Perkin-Elmer have produced a number of parallel image processors based on cellular logic operations, the latest called the PHP.[7] Chapter 10 of Preston and Duff offers a good survey of machines developed in the past two decades that can utilize the inherent parallelism of cellular logic operations on two-dimensional images.

As with image processing, researchers studying CLO applications have an important need to develop algorithms in the interactive mode, with graphical displays of the arrays. Toffoli[8] points out the general need for high-speed computing hardware to support the interactive exploration of cellular automata applications. Toffoli has developed a "Cellular Automaton Machine" for this purpose. In the spirit of this project, we have addressed this problem from a somewhat different approach than that taken for Toffoli's machine, Preston's PHP, or other recently developed processors.[7] Our goal is to achieve the least expensive CLO processor that still satisfies some minimum requirements of usefulness for a researcher.

By "least expensive," we mean we settled on a processor that would cost less than $500 and would be completely supported by a stand-alone IBM Personal Computer with a graphics board. We also required that the entire processor would fit on a single board in the PC. This cost goal easily brings the machine within the realm of any interested graduate students or independent researchers wishing to explore some cellular automaton ideas in the interactive graphics mode in their offices or homes. Such a computing task is not of wide commercial appeal; however, it is of great interest to a limited class of users.

We achieved high performance for very small cost by

- making the architecture narrowly focused and application specific,
- pipelining to obtain a reasonable level of parallelism, and
- using custom 3-micron chips for implementation.

Here we present the details of the processor along with its resulting performance.

Processor capability and architecture

A critical factor in realizing a practical application-specific processor is the early definition of a very specific form for the computation that is to be performed. The introduction of a lot of generality and flexibility into the design at extra cost obviously defeats the basic idea. We defined the computation for the CLO processor by

- selecting an arbitrary but reasonable array size,
- restricting the states to two, and
- applying the constraint of supporting interactive graphics displays. The processing speed needed for this constraint is difficult to achieve with software implemented on a general-purpose machine.

Conway's single-bit Game of Life provides a rich, standardized problem domain for cellular-automata research—and continues to stimulate a lot of ideas. Although processors for multibit cells had already been developed,[8] we felt that the two-state automata will continue to be useful for exploring new ideas for quite a while. We thus settled on a single-bit processor performing neighborhood logic operations on 256×256 arrays to greatly reduce the scope. The operation is based on the 3×3 neighborhood (that is, square tessellations). This neighborhood allows rules in the so-called von Neumann and Moore neighborhoods as well as hexagonal tessellation as a special case.

Enough processor speed and display speed to support interactive computing is much more important to the usefulness of a cellular automaton processor than the restriction to binary arrays. Toffoli's discus-

sion of this point suggests that something close to video rates (30 array transitions/second) is a useful display speed in this regard so that researchers can watch a full 2D array graphically "evolve under our eyes." This capability for a 256 × 256, single-bit array is just about satisfied by DMA transfers between the auxiliary cellular logic processor and the host PC memory. These transfers take about 50 milliseconds for the full array on the IBM PC.

For many applications, such as a sophisticated pattern-recognition algorithm, hundreds of CLO steps may be needed before a meaningful new state emerges. Consequently, we felt it was necessary to have the additional capability to compute transitions faster than video rate and then intervene periodically to update the graphics display. We achieved such a processor throughput in a simple architecture with a 12-transformation pipeline structure.

In this structure a cellular logic transformation on the rasterized array is performed at each sequential stage by a single, independently programmable, custom VLSI chip. Twelve such chips fit nicely on the processor board along with DMA servicing logic, and a 4-MHz clock is sufficient to allow the bit serial chip pipeline to keep up with the DMA byte transfers. The DMA bytes are serialized, and the bit stream snakes through the pipeline with the output of each (identical) chip being the input for the next. The host computer halts during the pipeline operation to maximize the DMA rate and is revived by an interrupt generated at the completion of the process.

This architecture allows 12 sequential state transitions of the complete 256 × 256 array in about 50 milliseconds, with the final state emerging from the pipeline via DMAs back to the computer memory or graphics memory. The parallelism introduced by pipelining 12 successive transformations helps keep the DMA transfers from becoming a computational bottleneck. The resulting effective processing rate of 240 array transformations per second can of course only be realized in a sustained mode for a problem in which the neighborhood rules for each of the 12 sequential transformations do not change from cycle to cycle. In cases requiring rule changes, the software intervention to load new transformation tables between processor calls adds a small additional delay.

The obvious key element in the processor is the custom chip that performs the cellular logic operations. This chip was developed as a class project in a VLSI course at Hopkins last year and was successfully fabricated through the MOSIS system at USC. MOSIS is a fast-turnaround chip fabrication service, initiated by the US Department of Defense Advanced Research Projects Agency and operated by the Information Sciences Institute at the University of Southern California. This service has greatly reduced the fabrication costs for small-volume projects by introducing process standardization and the multi-project wafer concept. In reasonable quantities the chips could be manufactured for less than $20 apiece.

CLO chip architecture

We designed our chip to perform, in a pipeline, a programmable square tessellation CLO transformation on a rasterized, single-bit array via a 64-byte lookup table to specify the Boolean mapping. The architecture is based on the original ideas presented by Golay[9] and their later implementations by the Perkin-Elmer group, although we have implemented the more-general square rather than the hexagonal neighborhood.

In addition, the chip keeps a count of the number of bits that are changed during the complete transformation. This count is an important piece of information for automating things such as image processing. For example, to perform skeletonization on an image that is an important preliminary step in pattern recognition, we apply the appropriate transformation repeatedly until the skeleton is stable. This stability is indicated by zero bits changing when the process is complete. The "bits-changed" counter for each chip is easily accessible as read-only memory locations mapped into the PC memory space.

The on-chip lookup table is a 512-bit static RAM accessible to the user as 64-byte read/write tables in the host PC memory space. During pipeline operation, as each bit is processed, its 3 × 3 neighborhood is used to create a 9-bit address into the lookup table. The high- and low-order bytes of the bits-changed counter are readable as table entries 65 and 66. Once a complete array has been processed, the bits-changed counter holds its value until it is reset by an external Sync signal that readies the chip for the start of a new array.

We used the standard approach of having the input array serially clocked into a 515-bit shift register (two complete rows plus three bits). The 9-bit lookup address is formed by judiciously selected cells along its length. The transformed bit is output directly into the next chip in the pipeline. A position (row/column) counter counts each input bit to keep track of array edges and the state of the pipeline so that various required control signals can be generated. One of these is a first-bit-out output signal that becomes the Sync (or start) signal for the next chip.

Two border-control chip inputs specify how to treat bits on the edges of the array. Four options are available: do nothing, invert the bit, make it a zero, or make it a one; the action is applied to the incoming bit stream. (If the do nothing option is selected, the architecture essentially wraps the array around the vertical axis to form a spiral of the horizontal rows.)

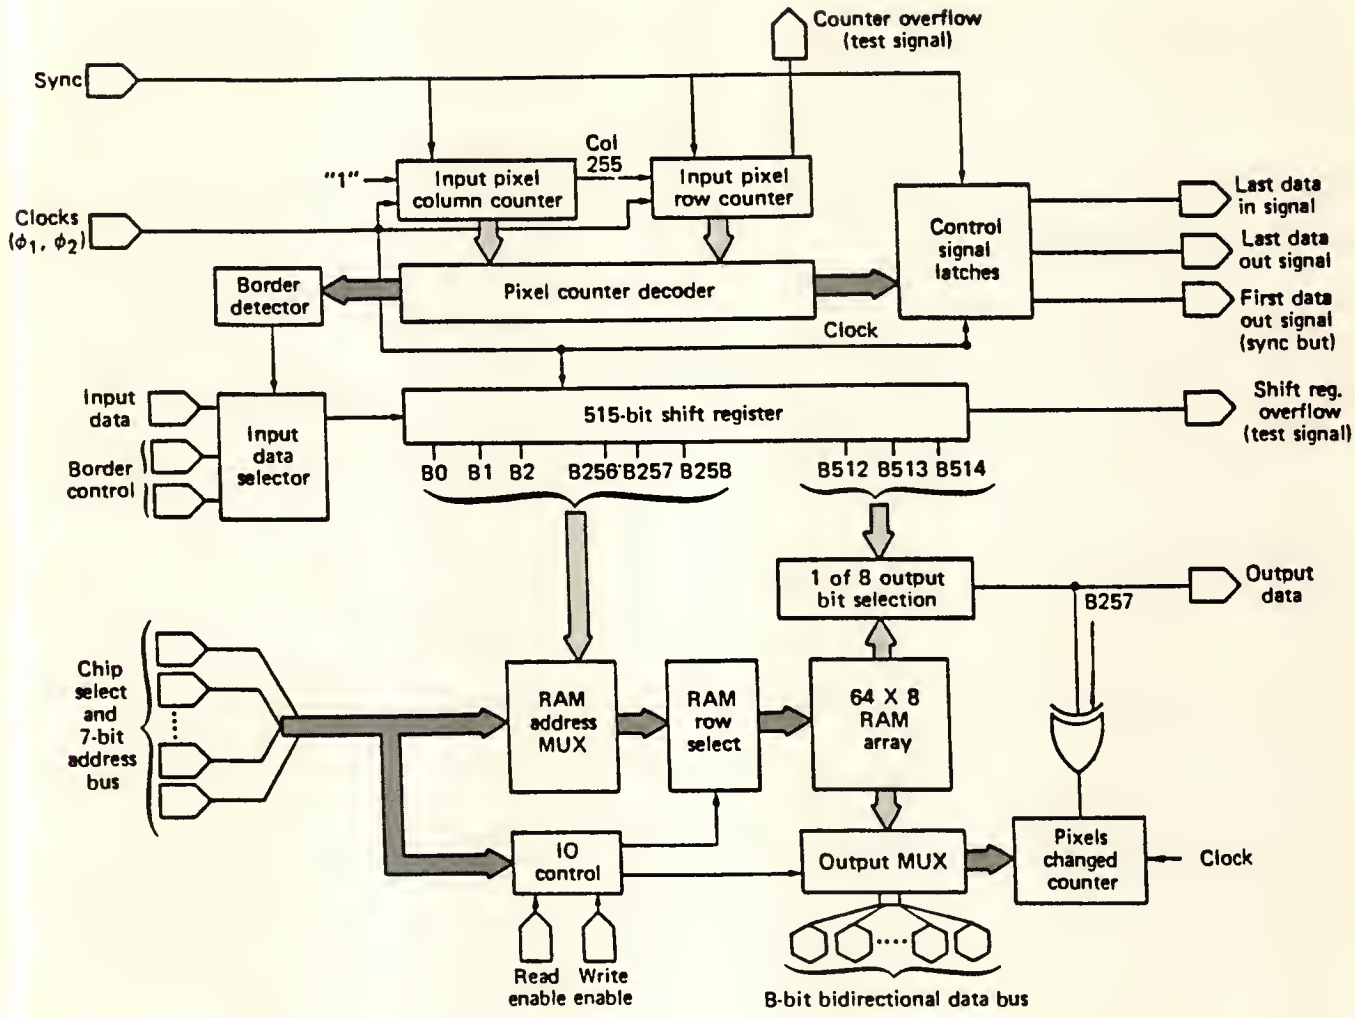


Figure 1. CLO chip architecture.

The chip architecture is shown in Figure 1. It is an NMOS design using one of the MOSIS standard 40-pin frames. The functional chips we received worked up to 10 MHz, which exceeded our requirements for an IBM PC coprocessor.

Architecture of processor board

The processor board contains: (1) the CLO chips forming the transformation pipeline, (2) logic to generate the two-phase clock that drives the chips from the 14.769-MHz system clock, (3) logic to handle the DMA byte transfer onto and off the board, and (4) address decoding for access to control registers and the chip lookup tables. Figure 2 displays the board architecture. An on-board DIP switch sets the base of the 2K address space of the processor.

Two channels of the PC's 8237 DMA controller are used in the single transfer mode to handle the processor I/O, with the output bytes lagging the input bytes by the total length of the pipeline (12 × 515 bits). Thus, the returned array can be targeted to the same memory block as the source array, if desired. This mode is handy for looping the pipeline since the DMA continuously resets itself automatically after 8192 transfers, and no software intervention is required in the loop to reset the DMA base addresses.

Parallel-serial interfaces are used at the input to the first chip and at the output of the last chip to interface to the DMA register on the board. After the software sends a Sync signal to the board to start the pipeline, the CPU is halted, and an Interrupt 7 request is issued by the CLO board to restart the CPU after the last DMA output. In this manner, only the

memory refresh cycle is active and can interfere with the DMA transfers, allowing them to proceed with maximum speed since this is the limiting factor on pipeline speed.

Benchmarking the performance

To benchmark the processor's power in image processing, we applied the processor to the Abingdon Cross benchmark case shown in Figure 3, which is also a nice example of some of the power of cellular logic operations. Preston at Carnegie Mellon proposed this benchmark as a processing standard. It has been used to formally test a number of existing and hypothetical parallel machines, with the results reported in Uhr et al.[10] The test involves the application of a series of standard, well-defined CLO transformations to first filter the image for noise, sharpen the feature edges, and then skeletonize the cross. The final result is shown in Figure 4.

Noise filtering by CLO transformations can be understood in terms of a simple example: the "grow" and "shrink" operations on a blob of 1's in a background of 0's. We perform these operations by specifying neighborhood rules that apply only to elements on the edge of a blob and have the effect of either eating inward by one cell or expanding outward by one cell. If a blob is large enough, a number of shrink steps followed by an equal number of grow steps will have no effect. However, a small blob of noise, say an isolated 1-cell, will disappear during the shrink steps and thus be eliminated.

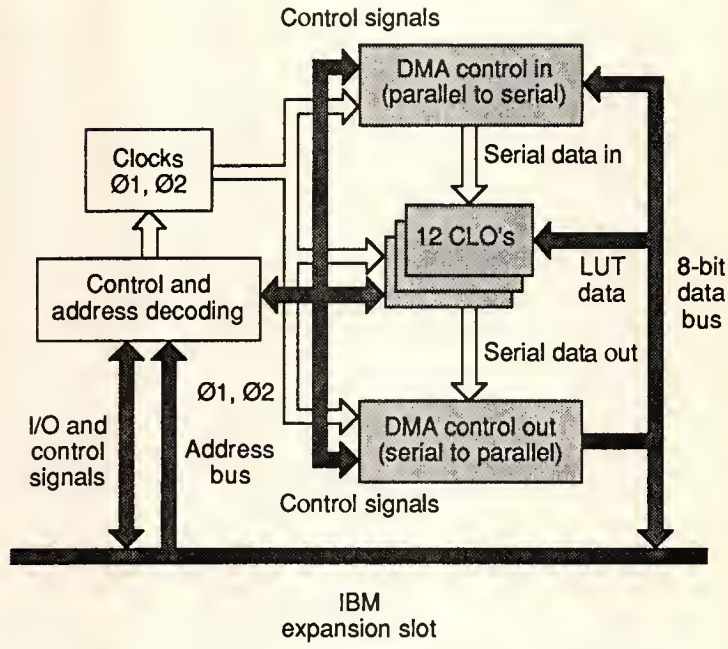


Figure 2. Processor board architecture.

Skeletonization uses a transformation similar to the shrink operation to erode the edges of any solid object inward symmetrically until a single line (or backbone) is produced.

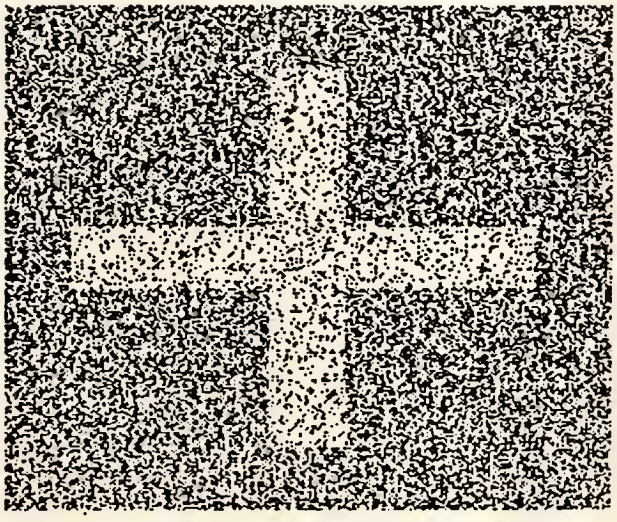


Figure 3. The Abingdon Cross benchmark image.

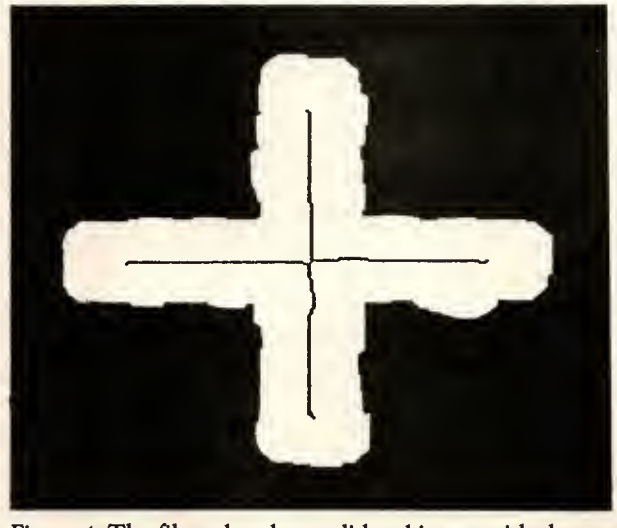


Figure 4. The filtered and consolidated image with the final skeleton shown superimposed.

Table 1.
Some systems benchmarked with the Abingdon Cross.[10]

Machine	Affiliation	Quality factor
CLIP4	University College London	7273
CYTO II	Environmental Research Institute of Michigan	39#
CYTO III	Environmental Research Institute of Michigan	102#
diff4	Coulter Electronics	17857 +#
DAP	International Comp. Ltd.	64000 +
DIP	Univ. Delft	46
FLIP	FIM/Karlsruhe	75
IP8500	Gould	†
IP8500	ETH/Zurich	78
Mach.Vision 2000	Machine Vision International	1924#
MPP	NASA Goddard	26283 +
PHP	Carnegie Mellon	7/31#
PICAP	Linkoping University	26
POP II	Royal Holloway College	2
System 75	International Imaging System	†
TAS	Leitz Wetzlar	533
TOSPICS II	Toshiba	206
TRAPIX 5500	Recog. Concepts, Inc.	40
VAP	University Berne	49
VICOM 10	Vicom	18
VICOM/CYTO	Vicom	6826 +
WARP	Carnegie Mellon	65*
CLO coprocessor	Johns Hopkins	175**

\# Algorithm used employs a priori information on width of cross.
\+ Estimated quality factor; benchmark not actually run.
† Investigator tried benchmark but could not obtain skeleton.
* Hypothetical system; benchmark estimated.
** Not included in the original Uhr et al.[10] table but added here for comparison.

The Abingdon Cross benchmark quality factor is defined to be the array size being processed divided by the clock-on-the-wall processing time for completion. The processing time should include all software intervention required to control the processor and its I/O. In our case the quality factor was about 175, based on an array size of 256 and a processing time of 1.46 seconds.

Table 1 shows some selected benchmark results from Uhr et al., all of which are for machines larger and more expensive than the CLO processor reported here. It is clear that the processor compares favorably to most larger machines for the computational task for which it was designed. For this project, the important question about performance deals with the comparison to a software implementation on a general-purpose computer. Performing 240 complete CLO array transformations per second on a 256 × 256 array requires approximately 15.7×10^7 individual element transformations per second. For implementation in software, if one assumes that only five CPU operations are involved for each individual transformation in the inner loop, this rate is equivalent to an effective processing rate of about 75 million operations per second. This performance at least equals the performance of most of the older vector mainframes on a well-vectorized problem. As a specific comparison to a software implementation, a C-language program on a VAX 11/780 produced about 0.9 array transformations per second of CPU time for this size array.

We feel a key element in the usefulness of the application-specific approach is developing the hardware as a coprocessor hosted by a personal computer. This approach is the fastest and most cost-effective way to achieve a stand-alone interactive processor. The host solves all the problems associated with interacting with the processor.

Another important element is that the processor is strongly decoupled from the software on the host machine by the use of simple block DMA transfers for processor data input and output. One benefit of this design decision is that the software required to support the processor is very uncomplicated. The few control signals needed (reset, border control, and start) are supplied by the user programs via memory writes to some absolute locations in the host computer's address space. The lookup tables for programming the neighborhood rules are handled the same way. High-level-language routines in the user's choice of language handle these functions easily.

Another advantage of the simple interface between host software and the processor is that maximum flexibility is achieved in the use of the processor. We discovered several related, but unanticipated, cellular automaton problems for which the coprocessor was able to be used. All of these resulted from the restriction of the processor to a very straightforward computing task and from the simple DMA interface used to move the data through the processor. The host machine's software assumes all the burden of managing the data's source and destination.

One example of such an unanticipated use is the case of a one-dimensional automaton array, which is an interesting computation in its own right.[6] One-dimensional cases can be handled by letting the initial conditions become the first row of a zero 2D array and setting the target base address of the DMA controller to lag the source address by one row (32 bytes). By scrolling the source address row by row

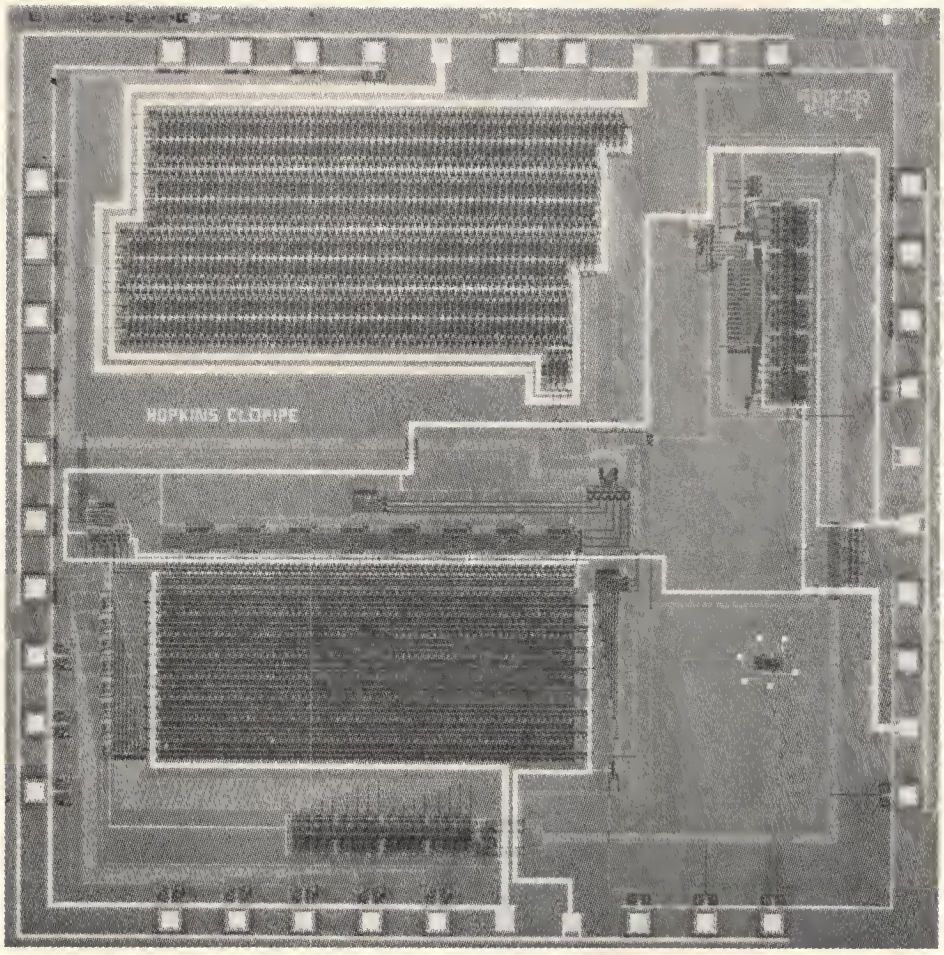


Figure 5. CLO chip laid out and simulated by students in the second semester of a VLSI design course, using the "university consortium" Unix-based CAD tools. The two major structures on the chip are the 64-byte RAM lookup table and a serpentine 515-bit shift register. The 7000 NMOS gates used in the design fit into a MOSIS-standard, 40-pin frame with obvious room to spare.

(modulo 256), one can graphically display a time history by rows, which is the usual method.[6] The software intervention required to reset the DMA base addresses each time the pipeline is called adds a small overhead in this case.

Another example is the class of transformations that are second order in time and satisfy time reversibility.[11] These transformations require that the transformation to a new state at $t+1$ be based on the two previous states at t and $t-1$. This requirement was handled by cycling the DMA source and target arrays through three separate arrays, thereby saving two predecessors of each new state. Through software intervention between pipeline calls, the correction for the earliest predecessor can be made. For two-state automata, such a procedure amounts to a simple exclusive OR of the two arrays, which can be handled by assembly code in 28 ms. The biggest throughput reduction in this case is achieved by the fact that only one transformation at a time can be handled by the pipeline, with the remaining 11 loaded with the identity transformation. Even so, we were able to watch graphics displays of the time-reversal rule referred to as Q2R in Vichniac[11] at a rate of 16 steps per second, which is definitely an interesting experience.

We estimate the total development time for the CLO chip (Figure 5) and the processor to be about six man-months. We felt this was the disappointing aspect of the project experience. One man-month seems more appropriate for the claim of nearly equivalent efforts in the hardware and the software. Clearly we have not quite reached our goal of truly fast-turnaround coprocessors, however the development time was short enough to give encouraging signals. Things are sure to improve in this regard; for example, we used traditional CAD tools for the chip layout rather than a silicon compiler that should reduce chip design time considerably. Also, the testing time required for the processor greatly exceeded the typical requirements for debugging equivalent software. More automated hardware testing techniques will surely improve this aspect within a few years.

On the other hand, for a computational task that will be run in production mode for a long period, a six-man-month development is a reasonable price for such a clear improvement in performance and throughput. For a certain class of problems in which the scope of the computing task is small and performance rather than generality is needed, the approach of using an application-specific coprocessor appears

preferable to implementation on a general-purpose machine. Even on the newly emerging highly parallel machines, the software effort involved might well be spent in producing a special-purpose processor. It is also not clear in many systems problems that are solved with an embedded microprocessor whether one might be replacing a hardware effort with an equivalent software effort and an attendant loss in performance.

Acknowledgments

Some of the architectural ideas for this processor were contributed by R.C. Moore and K. Strohbehn of the JHU/Applied Physics Laboratory. In addition the students in EE52.646 at JHU put some hard work into the CLO chip design during the spring semester of 1985. An NSF grant supported the chip fabrication.

References

1. K. Preston et al., "Basics of Cellular Logic with Applications in Medical Image Processing," *Proc. IEEE,* 1979, p. 826.
2. K. Preston, "Some Notes on Cellular Logic Operators," *IEEE Trans. Pattern Analysis and Machine Intelligence,* July 1981, p. 476.
3. M. Gardner, *Scientific American,* Oct. 1970, pp. 120-123.
4. L.S. Schulman and P.E. Seidan, "Statistical Mechanics of a Dynamical System Based on Conway's Game of Life," *J. Stat. Phys.,* 1978, p. 293.
5. C. Burks and D. Farmer, "Towards Modeling DNA Sequences as Automata," *Physica 10D,* 1984, pp. 157-167.
6. D. Farmer et al., *Cellular Automata—Proceedings of an Interdisciplinary Workshop, Los Alamos, N.M., 1983,* North-Holland Physics Publishing, Amsterdam, 1984, 245 pp.
7. K. Preston and M. Duff, *Modern Cellular Automata Theory and Applications,* Plenum Press, New York, 1984, 340 pp.
8. T. Toffoli, "A High Performance Cellular-Automaton Machine," *Physica 10D,* 1984, p. 195. (Also included in Reference 6.)
9. M.J.E. Golay, "Hexagonal Parallel Pattern Transformations," *IEEE Trans. Comput.,* 1969, pp. 733-740.
10. L. Uhr, K. Preston, S. Leviality, and J.B. Duff, *Evaluation of Multicomputers for Image Processing,* Academic Press, Orlando, Fla., 1986.
11. G. Vichniac, "Simulating Physics with Cellular Automata," *Physica 10D,* 1984, pg. 96. (Also included in Reference 6.)



Robert E. Jenkins developed the CLO chip as part of a course in VLSI design that he introduced as a visiting professor at Johns Hopkins Homewood campus in 1985. He is a lecturer in the Hopkins School of Engineering and a principal staff engineer at the Hopkins Applied Physics Lab. With this lab he has over 20 years of experience in computing systems applications and is currently leading the Space Department IR&D program.

A member of the IEEE, Jenkins received his BS in engineering and his MS in physics from the University of Maryland.



D. Gilbert Lee, Jr., performed the major portion of the processor board design and automated chip testing as part of an investigation of fast-turnaround, computer-aided design techniques. A senior staff engineer at the Hopkins Applied Physics Lab, he has experience in spacecraft digital subsystem development and is currently playing a leading role in applications of CAD tools.

A member of the IEEE, Lee received his BS in biomedical and electrical engineering from Duke University and his MS from Marquette University.

Questions concerning this article can be directed to Robert E. Jenkins, Johns Hopkins University, Applied Physics Laboratory, Johns Hopkins Road, Laurel, MD 20707.

Reader Interest Survey

Indicate your interest in this article by circling the appropriate number on the Reader Interest Card.

Low 156 Medium 157 High 158

Software Design for Real-Time Multiprocessor VMEbus Systems

Walter S. Heath

Taming real-time systems can be a challenge. You'll gain flexibility and cut debugging time by adding additional processors and the right software.

One can get into a lively discussion on the precise meaning of the term *real time*. Those of us who design and build real-time systems understand perfectly the meaning of the term within the context of our own work. It is quite another matter to come to a consensus on a general or global definition. Indeed, even within the industry we have no agreement on the correct spelling of the term (that is, real time versus realtime)!

Perhaps it is appropriate that the term used to describe these systems cannot be defined precisely, since a primary characteristic of real-time systems is that they must deal with uncertainty. Although a precise, quantitative definition of the term does not seem to be attainable, these systems do possess some common qualitative attributes.

Generally speaking, real-time systems are *reactive* in the sense that they are driven by external events. External equipment requests (or demands) the attention of the system, and the system must respond by performing a service within some prescribed response time. The interrupt mechanism is used to signal a request for attention, and so these systems are also classified as *interrupt driven*.

In this environment the software program must be designed to cope with uncertainty. That is, it must be capable of providing services to several external events that are occurring asynchronously and to do so within the response-time constraints of each external device individually. The program is therefore nondeterministic in that it is not possible to predict exactly what it will be doing a given number of clock cycles after initiation.

One has some limited control over the order in which services are provided by specifying interrupt and task priorities. Designers may assign highest priority to the interrupts from the most time-critical external devices and thereby ensure that they will be serviced as fast as the computer system is capable of responding. Similarly, they may assign scheduling priorities to application tasks. Beyond this, the designers must base their designs on a statistical average or worst-case analysis of system-response requirements. Since this analysis is, by nature, imprecise, the success of the design must rely heavily on the previous experience of the designers involved.

A discussion of the definition of the term *real time* often degenerates into a debate over what is meant by *immediate response*. This quantitative side of the definition is the most context dependent. For purposes of this article I define the phrase response time as the time interval between the occurrence of an interrupt by an external device and the first I/O operation performed by the computer in response to that interrupt. Systems implemented in custom hardware (perhaps controlled by microcode) may have response times in microseconds or less. On the other

0272-1732/87/1000-0071$01.00 © 1987 IEEE

hand, several minicomputer systems running large operating systems have a real-time mode. Response times here are usually in milliseconds. At the outer extreme are systems that respond in human response times (that is, fractions of seconds).

It is entirely possible for a single system to require different response times at different stages in the system. An example would be a system that must gather large quantities of raw data at high speed, process the data, and present results to an operator such that the operator has time to evaluate and possibly alter system operation in a timely manner. An air traffic control radar system is a good specific example. Here large quantities of A/D samples might be collected by high-speed, dedicated hardware under the control of a slower but more versatile microprocessor system. The data processing and display operations can require the computational power of a larger mini- or mainframe computer.

In this type of system the processing requirements change as data progresses through the system, beginning with a high-speed, dedicated (that is, dumb) front end and ending with the slowest but (presumably) most intelligent "processor"—the operator. The intermediate microprocessor stage may be neither fast enough to gather the raw data nor sufficiently powerful to perform data analysis. But it may possess adequate time response and versatility to control data collection and equipment configuration. As such, it serves as the central "facilitator" of operations by directing the operation of the dedicated front end and delivering blocks of data to the back-end computer.

In this article I discuss software design and development topics for microprocessor-based systems using multiple processors. Such systems currently have response times in the tens to hundreds of microseconds. As such, they may be capable of providing the entire processing component of a system, or they may serve as a stage in a more demanding system, as described earlier. I discuss system design considerations only with respect to their influence on the software. Details of specific software components appear in the accompanying boxes.

Serial versus parallel operation

The need to be capable of responding to several asynchronous events at essentially the same time has been a challenge for the traditional von Neumann single instruction stream, single data stream (SISD) or single-thread computer. The design requirement is inherently parallel while the computer is inherently serial. The traditional solution has been to provide a serial-to-parallel "transformer" in the form of a system executive that makes the computer appear to the outside world to be operating in parallel. The executive does this by time-sharing the computer's central processing unit between multiple tasks or processes. This so-called multiprocess or multitask operation is accomplished either by providing each process with a slice of CPU time, based on some priority/round-robin scheduling algorithm, or by allowing application tasks and interrupt handlers to dynamically schedule tasks to be run by a priority task scheduler. While this approach has been very successful and has been used to implement many real-time systems over the past several decades, it has some serious shortcomings, both technical and procedural.

Multitasking by a single CPU requires system overhead to transfer control from task to task. The entire state (registers, flags, and so on) must be saved for the task being suspended, the task scheduler must find another task to run, and the state of that task must be restored. This "context switch" can take tens to hundreds of microseconds.

Another problem with SISD systems is that, as new tasks are added to the system or as existing tasks are expanded, the performance of the entire system progressively degrades. This aspect has made designing such systems a risky enterprise, since one is not able to determine overall system performance until all components are functioning in their final forms. If the final system operates too slowly, the implementers must either produce more efficient code, remove less essential components, or purchase a faster version of the computer (if one is available).

The SISD approach also has some procedural difficulties. Since all software components must run together, a considerable coordination problem usually arises during the programming phase. For example, when parameters in a common data area must be changed, all programs that access that common area must be recompiled. It is also difficult to perform system integration, since the entire program must usually be run to test each component. While these problems can and have been surmounted in the past, current technology provides another approach that largely avoids them and provides additional capabilities as well.

The multiprocessor approach

A multiprocessor architecture matches real-time design requirements more closely. Multiple events can be serviced by multiple processors running in parallel. These systems are classified as multiple instruction stream, multiple data stream (MIMD) or multithread computers. Since less task switching takes place in each processor, executive overhead is reduced. The computational load is also distributed among several processors, thus relaxing individual processor execution time constraints.

The software in multiprocessor systems is often partitioned functionally; processors can be dedicated to performing specific parts of the job. An important

by-product of partitioning is that each individual processor program is simpler, is often written by one person, and can be tested independently. Thus processor functional partitioning leads naturally to the partitioning of software development responsibilities among personnel. Since programs can be written and tested independently, these efforts can proceed in parallel.

In a multiprocessor system a formal means for interprocessor communication must be established. Important system considerations in choosing the communication medium are:

1) hardware and software overhead to access the medium and
2) the data transfer rate after access is granted.

The MIMD class of computers can be further subdivided at this point into tightly coupled and loosely coupled subclasses. For systems in which high volumes of data must be shared between multiple processors with minimum latency, processors must be tightly coupled either by means of direct data paths or via a common multiport memory. For systems with less severe communications requirements a less tightly coupled and more flexible arrangement can be used.

Several media and communication protocols have been standardized for coupling processors in a loosely coupled configuration. For systems in which some communication speed may be traded for design flexibility, the backplane bus technologies are appropriate. Of these, the VMEbus has gained favor for many real-time applications.[1-4]

From the system designer's point of view an important feature of the VMEbus is its nonproprietary, open architecture. As a result, many suppliers are currently producing a wide spectrum of single-board computer (SBC) and I/O device interface boards. A short time after a new data or signal processing chip is introduced, it appears on a VMEbus board. It is therefore possible to configure the computer to match system requirements, rather than the other way around as was done in the past.

The VMEbus provides a high-speed, common interconnection between multiple SBCs and I/O devices. VLSI interface chips handle bus access protocol, so these operations are largely transparent to the software. The bus provides flexible priority interrupt and priority bus access request/grant mechanisms. I/O device ports and SBC dual-ported memory can be mapped into a single linear address space. One SBC can access the memory of another by simply generating addresses in its memory range; the SBC accesses I/O device ports as memory locations. Finally, blocks of data can be transferred across the bus at high speed using DMA devices, so that, on the average, the percentage of bus bandwidth being used can be kept low.

Passing Message Pointers

A previous *IEEE Micro* article describes a set of functions for passing messages among application tasks and interrupt handlers using queues.[5] For some applications it is preferable to pass pointers to messages, rather than the messages themselves. This is particularly true when the messages are large. Passing a pointer takes considerably less time than passing an entire long message. Since pointers are fixed in length, the queues have fixed-size entries, and so the queue access functions are simpler.

A buffer pointer queue consists of a queue header and an array for storing pointers to buffers. The header has the following structure:

```
struct pque{          /*buffer ptr queue header  */
      char sema      ;/*access semaphore         */
      char full      ;/*queue-full flag           */
      short head     ;/*head pointer (index)      */
      short tail     ;/*tail pointer (index)      */
      short lngth    ;/*queue length (# of ptrs)  */
      short count    ;/*count of msgs in queue    */
      char task      ;/*waiting task ID           */
      long *pbuf     ;/*ptr to ptr queue array    */
}                    ;
```

A queue's *head* pointer (index into the pointer array) points to the next buffer pointer to be removed. A queue's *tail* pointer points to the next open entry in the queue. Note that these pointers will end up pointing to the same entry when the queue is either full or empty! The Full flag is used to resolve this ambiguity. The pointer queue's length is specified by *lngth*. This flag is used to determine when to *wrap* a head or tail pointer (that is, move it from the end to the start of the pointer array).

Actual access to a queue is performed by primitive functions Putbp() and Getbp() (put/get a buffer pointer), so the user does not need to manipulate parameters in the queue header directly. If these functions are unsuccessful (the queue is full or empty), they return -1. Various higher level functions that supply useful additional services call these primitive functions. For example, functions Putbpwt() and Getbpwt() (put/get a pointer or wait) use a queue header variable *task* to determine if another task attempted to access the queue and was unsuccessful and therefore suspended action (called Sleep().) If a number other than minus one is present in *task*, it is interpreted as the ID of a suspended task and that task is awakened.

Thus, if Putbpwt() places a pointer in a queue and finds that another task has called Getbpwt() in an attempt to get a pointer and found none was available, it will wake the suspended task. Since a pointer is now available, that task can proceed. Similarly, if Putbpwt() attempts to place a pointer in a full queue, it will suspend. Then, when Getbpwt()

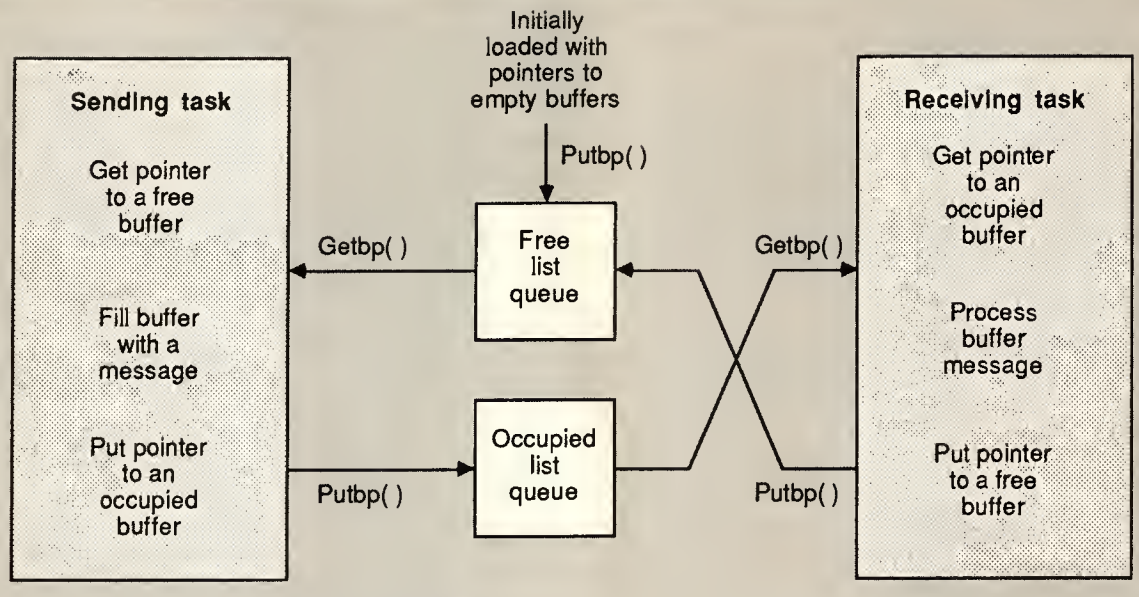


Figure A. Procedure for using buffer pointer queues.

removes a pointer, it will wake the waiting task. These two functions therefore provide a mechanism that allows data flow to control task execution.

Obviously, the pointer to the array of pointers *pbuf* and all of the other header variables must be initialized before they can be used. An initialization function performs this duty as part of the executive's start-up procedure.

The procedure for using buffer pointers in an application is diagrammed in Figure A. For a particular intertask (or task-interrupt handler) path, a set of two buffer pointer queues is defined, one a *free list* and the other an *occupied list*. The free list queue is initially loaded with pointers to fixed-size message buffers, using Putbp(). Then, when a task needs to send a message, it calls Getbp() (or a higher level calling function) to remove a pointer to a free buffer from the free list queue, loads the buffer with data, and then places the buffer's pointer in the occupied list queue, using Putbp(). The receiving task, when it needs a message, calls Getbp() to remove a message pointer from the occupied list, processes the message, and returns the pointer to the free list, using Putbp().

Note that the pointer queue functions are not concerned with the contents of the messages passed. Obviously the application program must be aware of message format and size. If message size is variable, it can be sent as part of the message. It is also clear that a set of fixed-size buffers that are as large as the largest message sent must be defined for each message path. The maximum number of messages that can be queued at any one time is also fixed. Finally, the execution time required to pass access to a message is fixed, so execution time is independent of message size.

By contrast, when messages themselves are passed in the queues, the number of messages that can be queued depends on the number that can be placed in the queue array. A large number of small messages can be queued or a smaller number of larger messages can be stored. In addition, the same queue space is used repeatedly as messages are passed to and from the queue. These operations therefore use memory space more efficiently. But the time required to pass a message is directly proportional to its length.

When passing message pointers between SBCs, a semaphore flag must be used to control access to the queue. Otherwise two SBCs might attempt to manipulate a queue header at the same time. The *sema* flag controls access. A special assembly language instruction must be used to manipulate this flag. For example, the 68000's *tas* (test and set) instruction indivisibly tests and sets the flag. The instruction returns the setting of the flag previous to instruction execution (in the processor's Z status bit). By convention, when sema is set, access to the queue is blocked. Thus, when tas is executed, it blocks access and returns the previous access condition. If the queue was not already blocked (sema was cleared), the calling SBC can then access the queue. If the queue was blocked, it must wait. When an SBC is finished with a queue, it must clear sema to allow access by the other SBC. An indivisible instruction is required to avoid the situation where both SBCs test the flag, find it cleared, set it, and then attempt to access the queue at the same time.

Functions Putbpr() and Getbpr() (put/get buffer pointer or retry) provide inter-SBC queue operations. These functions test the semaphore flag before calling Putbp() or Getbp(), respectively. A calling argument specifies the number of times the semaphore flag should be polled before failure is reported (-2 returned).

System design considerations

Any common bus system should be designed to minimize bus traffic. Although VMEbus data throughput is high, the bus is usually the only data channel between the SBCs and is therefore a potential bottleneck. Bus traffic can be reduced by passing message buffer pointers rather than the messages themselves. When the final destination for the message has been determined, the message can be passed at high speed using a DMA device. Wherever possible, I/O should be performed directly from the SBC board using on-board I/O devices. Several SBC boards contain serial and parallel ports and disk controllers that perform I/O either via the board's P2 connector or through the front panel. The use of these ports removes some of the traffic from the bus.

Another way to reduce system design risk is to simply design in computational margins at the beginning. Since SBCs are relatively inexpensive, it is not unreasonable to use more of them than the initial design requires so that each one operates below its potential capacity. Since smaller programs in more SBCs are easier to write and test than larger and fewer programs, the extra cost of additional SBCs will usually be offset by lower software production costs. In addition, since timing constraints in each SBC are relaxed, programs can be written in a high-level language; the associated execution time penalty is usually of little concern.

It is important that system integration and debug procedures be considered during the initial design phase. Individual SBC programs can be tested using the debugger in the SBC's PROM monitor or perhaps by using an in-circuit emulator. It may be necessary to write separate source and destination simulator programs and run them in adjacent SBCs to simulate bus traffic to and from the program being tested. These can be simple, menu-driven loop programs.

It is also important to take into account the types of data processing that will be performed. Common types of operations can then be grouped together in dedicated SBCs. For example, operations can often be partitioned into computation-intensive, communication-intensive, and I/O-intensive categories.

A typical design

Figure 1 depicts a candidate real-time VMEbus system. It contains a bus arbiter board and several SBC and I/O boards. A design might contain a central dispatcher SBC that collects message pointers from other SBCs, determines message types, and dispatches them to appropriate destination processors. Since this processor is the center of activity, it might also support data transfers for the system's most time-critical I/O devices. For example, it might control message traffic to and from a high-speed communications channel.

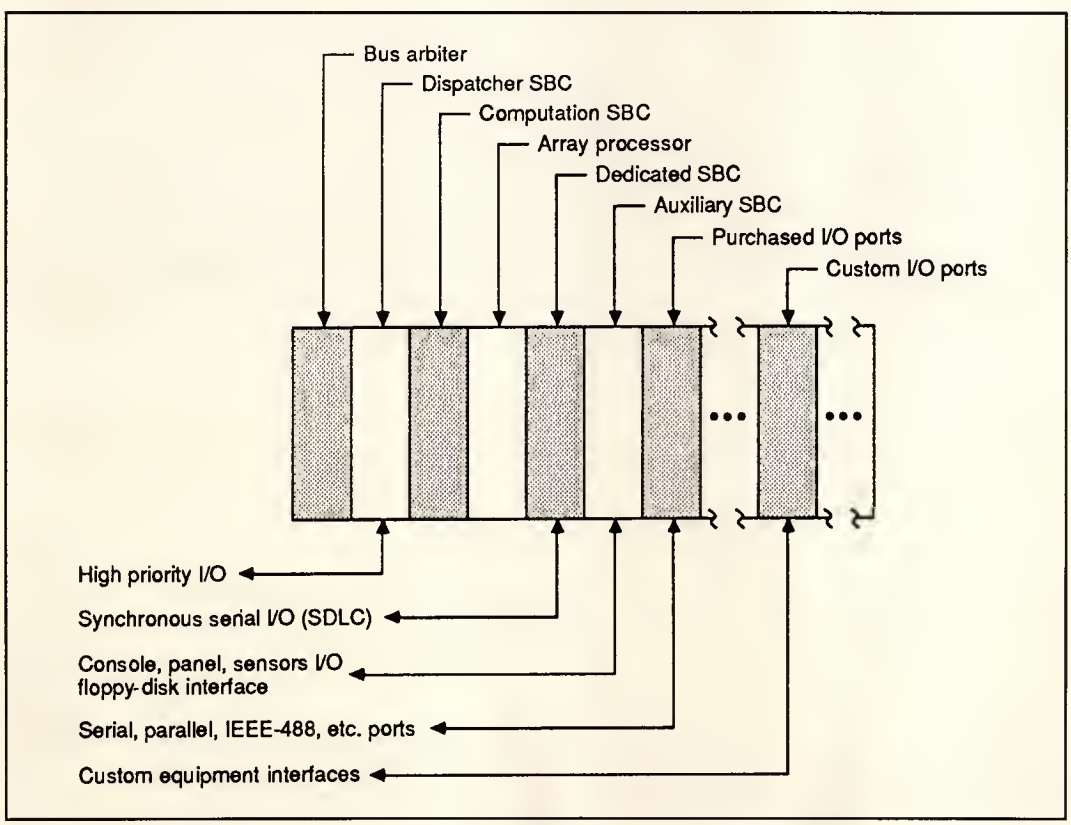


Figure 1. Typical real-time VMEbus system components.

Communicating Between SBCs

In a multiprocessor system designers must establish a formal procedure for communicating between processors. Since each SBC's dual-ported memory is separately mapped into the common address space of the system, it is possible to establish common data areas for communication. C language programming techniques that can be used to establish and access these commons follow.

A data structure type can be declared that contains all variables and data structures to be placed in the common area. An example structure appears in Figure B.

In this example I have established a set of buffer pointer queues between source and destination SBCs. The example also includes parameters for reporting status and error conditions (an auxiliary SBC might collect this data and report it on the operator's console). Finally, the Ready and Go flags used during program start-up appear (see main article).

In the source and destination SBC programs a pointer to a structure of type sbccom can be declared and initialized as follows:

```
#define SBCSTART  0x400000
struct sbccom *sp                ;
sp = (struct sbccom *)SBCSTART   ;
```

In this example the pointer is declared to point to absolute hex address 0x400000, which may be either an address of memory on the source or a destination SBC board (or on another board). Since sp points to this address and is declared as a pointer to structure sbccom, the sbccom data structure overlays a memory area starting at the pointer address. It is then possible for source and destination programs to access parameters in the sbccom structure using pointer sp. For example, the source status can be initialized to zero as follows:

```
sp->srcstat = 0                  ;
```

Similarly, message buffer pointer msgptr can be placed in queue osrcdst using function Putbpr() as follows:

```
putbpr(msgptr,&(sp->osrcdst),1,1)   ;
```

As noted in the box titled Passing Message Pointers, function Putbpr() uses a semaphore to control access to the queue. Since the VMEbus supports indivisible instruction execution (the 68000's tas instruction), it is valid to use the semaphore technique to control access to queues that are accessed by multiple SBCs.

This procedure for passing parameters and buffer pointers can be used to establish common areas between the various SBCs in the system. By using several common areas, less recompilation is necessary when parameters are added or changed. Structure declarations such as sbccom should be placed in separate include (.h) files in a common library so that the current version is always included when programs are recompiled.

Note that a process must poll a queue to gain access to it and then to determine whether a message pointer is present or whether room exists to store a pointer (the last two arguments in the Putbpr() call specify the number of times the function should poll before returning failure). Experience has shown that this is not a serious problem. If an SBC is busy processing a previously received message, it should probably not be interrupted; if the SBC is not busy, it might as well spend its time polling its queues. Queues that can be expected to require substantial polling should be placed in the memory of the SBC that will be doing the polling so that excessive polling will not occur across the bus.

```
struct sbccom{
    struct pque fsrcdst          ;/*free list queue header     */
    int fsdbuf [FSDLNGHT]        ;/*free list ptr buffer        */
    struct pque osrcdst          ;/*occupied list que header */
    int osdbuf [OSDLNGTH]        ;/*occupied list ptr buffer    */
    int srcstat                  ;/*source status               */
    short srcerr1                ;/*source error #1 flag        */
    short ready                  ;/*SBC-ready flag              */
    short go                     ;/*SBC-go flag                 */
}                                ;
```

Figure B. Sample data structure.

Most systems will also need an auxiliary SBC to handle lower speed interfaces, such as the operator console, control panel interfaces, and/or environmental sensors (temperature, position, vibration). One or more SBCs might be dedicated to performing computation-intensive operations. These processors might contain conventional CPUs with attached coprocessor floating-point math chips, or they may be array processors. In some designs SBCs might be dedicated to supporting specific I/O operations. To get maximum performance, certain high-speed I/O devices require the undivided attention of a separate processor (a high-speed, synchronous serial communication link).

A typical system will also contain several custom and/or purchased I/O boards. Boards for most standard I/O protocols are available. Special-purpose interfaces can be fabricated using wirewrap techniques. In all cases the software interfaces via memory-mapped registers and interrupts. Some boards include DMA chips to support high-speed I/O transfers.

The system executive

The software that runs in the SBCs can be divided into two general categories, system software and application functions. A system executive provides an environment within which application programs can be constructed and run. As with governments, a system executive should be as small as possible—consistent with the need to supply essential services. An executive should not waste machine cycles on unnecessary or inefficient operations. A cycle lost to system software is a double loss to the application since both a cycle of useful computation and a cycle of real time are lost. For high-speed systems the executive should therefore be designed to meet the specific needs of the application as closely as possible. It should also provide a convenient interface to application programs and should not be so complex as to be difficult for users to understand and use.

At minimum, an acceptable system executive should provide priority task scheduling, support interrupt processing, and provide formal mechanisms for intertask and interprocessor communications. But, some of the services provided by even this small set of functions may not be needed in some applications. For example, if system requirements dictate that each SBC in a multiprocessor system should run only one task, a task scheduler is not needed. In systems that must provide the highest possible performance, unnecessary executive components should not be included.

One must decide whether to purchase an executive or write one. A custom executive can be designed to meet the specific needs of a project and can be modified, if necessary, as application requirements change. If the source code for the executive is avail-

The System Executive and Preemptive Scheduling

The system executive design described in a previous *IEEE Micro* article provided priority, nonpreemptive task scheduling, and utility functions for passing entire messages using queues.[5] Since then, the executive has been extended to include functions to pass message pointers between tasks and between processors. Here, I briefly review the work and present some thoughts on preemptive scheduling.

Figure C presents a structural overview of the executive. It operates as follows. The user must define a Task Control Block (TCB) data structure and a stack for each application task. TCBs are linked in a list, with the last TCB linked to the first, to form a loop. The scheduler function Sched() cycles through TCBs looking for a task that has been awakened. The executive spends its time in this way if no tasks have been awakened.

An interrupt handler or another task can wake a task by calling function Wake(). When Sched() finds an awakened task, it calls Run(). Function Run() saves Sched()'s stack pointer (SP) and transfers control to an entry point in Sleep() (using assembly code). Function Sleep() gets the task's SP from its TCB and returns to the task. (The return address was pushed onto the task stack during a previous

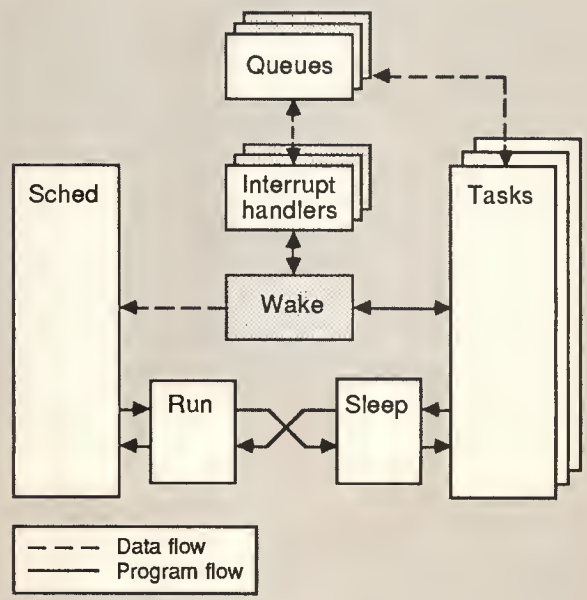


Figure C. System executive, structural overview.

call to Sleep().) The task loop performs user-supplied operations and eventually calls Sleep(). Sleep() transfers control back to Run(), and Run() returns to Sched(). The task's SP is saved, and Sched()'s SP is restored in this operation. Since Sched() always starts scanning TCBs from the top of the linked list, the order of TCBs on the list determines task priority.

When an interrupt occurs tor a high-priority device, the interrupt handler may need to cause the interrupted task to suspend action so that the data received by the handler can be processed by a higher priority task. That is, the lower priority task is preempted by the higher priority task. Note that this procedure is needed only in closed-loop feedback situations in which the received data must be processed immediately so that an output operation can be performed within some small time interval. In most real-time applications this level of performance is not needed and should not be used, since it imposes additional system overhead and application programming constraints. It also makes program debugging more difficult.

If a task is preempted when it is in the process of manipulating a data structure (for example, a queue header), that structure is left in an inconsistent state. If the preempting task then accesses the same structure and changes its contents, the preempted task has an inconsistent structure to deal with when it again runs. Clearly it will be necessary to disable preemption when certain "critical regions" of code are being executed. A similar situation arises for common functions that may be called by both tasks (C library functions, PROM utilities). The partial results from the call from the preempted task must be saved so that they are not overwritten by the call made by the preempting task. Functions that are programmed to cope with this situation are called reentrant.

To support preemptive scheduling, an executive must be capable of providing a mechanism that will cause the preempted task to suspend at the point of interruption. The executive must then reschedule the task so that the next time it is selected by the scheduler it will run from the point of interruption. That is, it must pause to allow the scheduler to run the higher priority task. Suspension can be accomplished by manipulating the data that is pushed onto the supervisor stack when the task is interrupted.

The only penalty in using nonpreemptive scheduling is that an interrupt handler may awaken a higher priority task and that task will not get a chance to run until the current task suspends. In actual systems this problem can usually be solved by inserting calls to a Pause() function at various points outside critical regions in the lower priority tasks.

Pause() simply calls Wake() to reschedule the calling task and then calls Sleep() to suspend the task. This procedure allows the scheduler to check to see if a higher priority task has been awakened and to run it if it has. When the task that called Pause() becomes highest priority, it will be resumed by a return from the Pause() call. Insertion of Pause() function calls is a fine-tuning exercise that is usually performed in the latter stages of system integration to improve overall system performance.

Another way to avoid preemptive scheduling in a system that must support fast response to some critical I/O interface is to simply devote a separate SBC to handling that device exclusively. In this approach the program running in the SBC is likely to consist of only one task. It is therefore unnecessary to preempt the task to start another. An interrupt handler can simply return without having to check to see if the running task should be preempted. This method provides the fastest possible interrupt response time.

I believe that preemptive scheduling adds unnecessary risk and complexity in all but a few designs and that ways can usually be found to avoid it in those few remaining situations. A nonpreemptive environment is more predictable. A task running in this environment maintains control of the CPU until it chooses to give up that control (except for interrupt servicing). Thus, a task can be assured that a non-reentrant function has been completed or that all items in a data structure have been updated before another task is allowed to access that function or structure. It is therefore unnecessary to suffer the overhead of locking and unlocking code in critical regions.

able, users will have a complete understanding of every detail of the SBC programs they are writing. This knowledge contributes considerably to software productivity.

I described an executive design that is appropriate for many VMEbus systems in a previous article.[5] This executive is designed to be written primarily in the C language and in assembly language, where necessary. It can be placed in a library in a development system and linked to application tasks and interrupt handlers. The original design and extensions are discussed in the accompanying boxes. This executive is "lean," and therefore fast, for an executive that is designed to be written in a high-level language.

An executive written exclusively in assembly language would be faster but would be more difficult for users to understand. In a particular application, if extreme speed were needed, certain critical regions could be coded exclusively in assembly language.

The executive design has been extended to support the passing of pointers to message buffers between tasks within a processor and between processors via the bus. Buffer pointer queues pass the pointers. It is important to avoid unnecessary transfers of blocks of data within processor programs and especially across the bus. Instead, a pointer is passed. The destination task can then examine the message header and select the next processing step. Once a decision is made on the ultimate destination for the message, the pointer can be used to make a single, high-speed block transfer. Since all dual-ported memory and I/O ports are mapped into one linear address space, it is easy to accomplish a block transfer either by means of a program loop or by a DMA device.

Initial program load and start-up

In the process of writing and debugging individual processor programs, it is convenient to be able to download memory images from the development system to the target VMEbus processor using a serial link (I discuss software development next.) But during system integration and in the final system it is more convenient to be able to automatically load programs from a floppy disk. Since each SBC's dual-ported memory is separately mapped into the common address space of the system, the SBC performing the autoload has little difficulty transferring programs to other SBCs over the bus. The PROM monitor for the autoload SBC usually must be extended to support disk operations.

The modified PROM monitor program should be capable of reading a command (script) file from the disk. This file should contain a list of names of memory image files to be loaded. It should then be capable of reading these files consecutively into the memories of the various SBCs. If the software is developed on a Unix (or PC) system, the PROM monitor can be modified so that files can be read from a Unix (or PC-DOS) directory on a floppy disk. A file will then have a standard header that will contain the program's starting address and size.

Once the programs are loaded, they can be started by means of front-panel toggle switches on the individual SBCs. Reprogramming the Abort (or other) front-panel switch to start the program (another PROM patch) accomplishes this task. Each SBC can then perform start-up initialization operations and set a Ready flag in a common memory region. When all SBCs have set their Ready flags, the auxiliary processor can set Go flags for each SBC in a sequence that is appropriate for the application. Each processor then proceeds to run its real-time program. (SBC suppliers please take note that it would be a great help to users if the PROM monitor additions mentioned here could be standardized and included in delivered PROMs.)

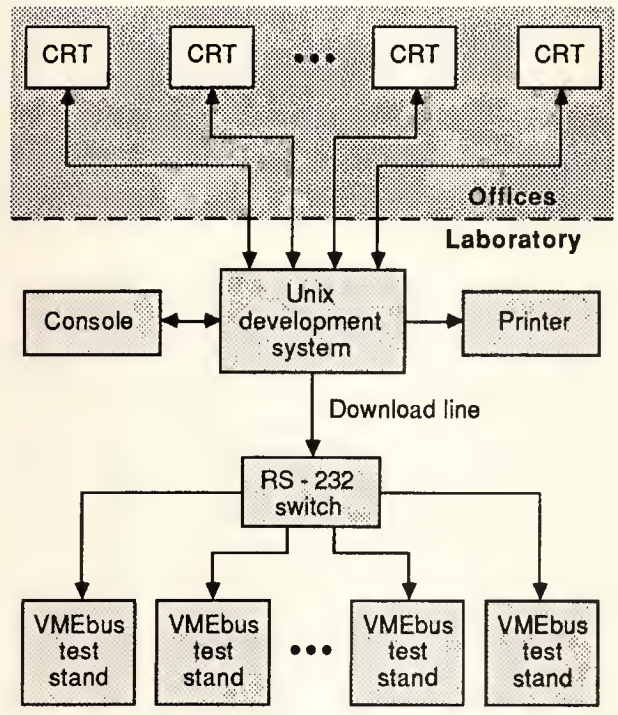


Figure 2. Software development configuration.

Software development

A good environment for developing processor programs is a small Unix system that is using the same CPU chip as the target SBC boards. The system's Unix compiler, assembler, and linker can then be used to write application programs. Otherwise a cross-compiler is needed. A Unix utility is also available for generating a memory load map, which is needed during debugging. Unlike some larger minicomputer systems with proprietary operating systems, a small Unix system can be supported by the people using it. There is no need to have a system administrator with special knowledge of the system to keep it running.

An efficient equipment configuration is shown in Figure 2. It consists of a Unix system with serial lines to user's offices for terminals, a separate VMEbus "test stand" for each user in a central laboratory location, and a serial port for downloading memory

images to the user VME systems. This download line should have a switch so that downloads can be directed to the various test stands. By having a separate test stand, each user can configure a system with the necessary SBC and I/O boards required to test software. In this way users are not delayed by waiting for access to test equipment.

At first glance the multiprocessor approach to designing real-time systems might appear to be overly complex and risky. It could be argued that a more traditional single-processor design would be safer. But experience has shown that the reverse is often true. Additional processors provide more design flexibility, more closely match the parallel nature of real-time systems, and provide more computational throughput. It is also easier to produce, debug, and integrate several simpler real-time programs than a single complex program. The inherent parallel nature of the system also means that both hardware and software development activities can proceed more smoothly in parallel. In short, multiprocessor systems reduce design risk, improve performance, and increase productivity.

Acknowledgments

This article is based on projects performed at Lincoln Laboratory, Lexington, Massachusetts. Over a dozen people contributed to these projects. The following are representative. Gary W. Ahlgren was the principal systems architect and source for solutions to complex technical problems. Tim Bowler, Marie A. Donnelly, Ellen M. Jervis, Anne M. Matlin, Linda Riehl, and the author produced software. Carol C. Martin provided valuable advice on the content and structure of this article.

This work was sponsored by the US Department of Defense. The views expressed are those of the author and do not reflect the official policy or position of the US Government.

References

1. *VMEbus Specification Manual, Rev. C.1,* Vita, 10229 North Scottsdale Road, Suite E, Scottsdale, AZ 85253; (602) 951-8866.
2. *VMEbus Systems* (magazine), 25875 Jefferson at Madison, St. Clair Shores, MI 48081; (313) 774-8180.
3. W. Fisher, "IEEE P1014—A Standard for the High-Performance VME Bus," *IEEE Micro,* Feb. 1985, pp. 31-41.
4. J. Hemenway, "Powerful VME Bus Features Ease High-Level μC Applications," *EDN,* Jan. 12, 1984, pp. 158-165.
5. W.S. Heath, "A System Executive for Real-Time Microcomputer Programs," *IEEE Micro,* June 1984, pp. 20-32.



Walter S. Heath is a computer systems designer with 20 years of experience in building real-time embedded computer systems. He has spent the last six years at Lincoln Laboratory on a contract basis building research prototype systems, including contributions to the TCAS Airborne Collision Avoidance System and a GPS navigation system. His previous experience includes the design and implementation of computer navigation, communication, and control systems including initial computer implementations of Loran and Omega radio navigation systems.

Heath received BSEE and MSEE degrees from the University of Michigan, where he specialized in electromagnetic field theory and computer technology.

Questions concerning this article can be directed to the author at 118 Fairhaven Road, Concord, MA 01742. Readers interested in further study of real-time system software can obtain the source code for the executive and example real-time application programs for 680X0 processors on a PC 5.25-inch floppy disk from Intr-Soft Co., PO Box 351, Bedford, MA 01730; (617) 369-6242.

Reader Interest Survey

Indicate your interest in this article by circling the appropriate number on the Reader Interest Card.

Low 168 Medium 169 High 170

1987 Index

IEEE Micro Vol. 7



This index covers all technical items that appeared in this periodical during 1987, and items from previous years that were commented upon or corrected in 1987.

The *Author Index* contains the primary entry for each item, listed under the first author's name, and cross-references from all coauthors. The *Subject Index* contains several entries for each item under appropriate subject headings, and subject cross-references.

It is always necessary to refer to the primary entry in the *Author Index* for the exact title, coauthors, and comments/corrections.

SUBJECT INDEX

A

Application-specific integrated circuits
application-specific coprocessor for high-speed cellular logic operations. *Jenkins, R. E., + , M-M Dec 87* 63-70
logic elements and symbol creation for designing with logic cell arrays. *Landry, Steve L., M-M Feb 87* 51–59

Arithmetic
approximating binary logarithm with integer arithmetic. *Männer, R., M-M Dec 87* 41-45
IMS T800 transputer scientific computer with improved floating-point and communication performance; design and architecture. *Homewood, Mark, + , M-M Oct 87* 10–26

B

Bibliographies
minicomputer-based and microcomputer-based FFT implementations for advanced measurement systems. *Van der Auweraer, H., + , M-M Feb 87* 39-49

Binary arithmetic; cf. Arithmetic

Book reviews
CD ROM 2: Optical Publishing (Ropiequet, S., et al., *Eds.*; 1987). *Mateosian, Richard, M-M Jun 87* 83–84
MC68851 Paged Memory Management Unit User's Manual (Motorola, Inc.; 1986). *Mateosian, Richard, M-M Apr 87* 84
Microprocessor-Based Design, A Comprehensive Guide to Effective Hardware Design (Slater, M.; 1987). *Mateosian, Richard, M-M Aug 87* 85–86
Operating Systems Design—The XINU Approach (Comer, D.; 1984). *Mateosian, Richard, M-M Oct 87* 90–91
The Art of Desktop Publishing, 2nd edn. (Bove, T.). *Mateosian, Richard, M-M Jun 87* 84
The Connection Machine (Hillis, W. D.; 1985). *Mateosian, Richard, M-M Apr 87* 84–85
The Design of the UNIX Operating System (Bach, M. J.; 1987). *Mateosian, Richard, M-M Oct 87* 91

C

CAD (computer-aided design); cf. Design automation

CD-ROMs
book review; CD ROM 2: Optical Publishing (Ropiequet, S., et al., *Eds.*; 1987). *Mateosian, Richard, M-M Jun 87* 83–84

Cellular logic
application-specific coprocessor for high-speed cellular logic operations. *Jenkins, R. E., + , M-M Dec 87* 63-70

Communication switching
low-cost nonstandard LAN for controlling family of small telephone exchanges. *Micheletti, Giancarlo, + , M-M Oct 87* 70–82

Computer-aided design; cf. Design automation

Computer architecture
capability-based microprocessor system architecture. *Corsini, Paolo, + , M-M Jun 87* 35–51
Clipper 32-bit microprocessor; system architecture, programming model, and memory management features. *Hunter, Colin B., M-M Aug 87* 6–26
emulating complex instruction set computer with reduced instruction set computer using GaAs circuits. *McNeley, Kevin J., + , M-M Feb 87* 60–71
European approaches for advanced architectures (special issue). *M-M Oct 87* 4–82
Japan's TRON project (special issue). *M-M Apr 87* 3–80
TRON project; architecture of TRON VLSI CPU. *Sakamura, Ken, M-M Apr 87* 17–31
TRON project, Japan's family of specifications for computer architectures, operating system kernels, and VLSI chips; overview. *Sakamura, Ken, M-M Apr 87* 8–14
80387 floating-point coprocessor architecture, interface protocols, and applications. *Perlmutter, David, + , M-M Aug 87* 42–57

Computer arithmetic; cf. Arithmetic

Computer fault tolerance
fault tolerance and reliability for process control computers; overview and examples of European computers. *Kirrmann, Hubert D., M-M Oct 87* 27–50

Computer industry
European cooperation in information technology industry; Esprit program, *McLauchlan, Derek J., M-M Oct 87* 6-9

Computer interfaces; cf. Microcomputer interfaces

Computer languages
DOOM (Decentralized Object-Oriented Machine) for executing programs in POOL (Parallel Object-Oriented Language). *Bronnenberg, Wim J. H. J., + , M-M Oct 87* 52–69
system performance modeling for complex VLSI; PAWS simulation language. *Iacobovici, Sorin, + , M-M Aug 87* 59–72

Computer languages; cf. Microcomputer languages

Computer networks; cf. Local area networks; Microcomputer networks

Computer operating systems; cf. Software, operating systems

+ Check author entry for coauthors

† Check author entry for subsequent corrections/comments

Computer performance
system performance modeling for complex VLSI; PAWS simulation language. *Iacobovici, Sorin, + , M-M Aug 87* 59–72
Computer reliability; cf. Computer fault tolerance
Computers; cf. Distributed computing; Microcomputers; Parallel processing; Supercomputers
Control systems; cf. Process control
Copyright protection
copyright protection for screens and interfaces for computer programs (MicroLaw). *Stern, Richard H., M-M Jun 87* 81–82
NEC vs. Intel; implications for microcode, instruction sets, and compatibility (MicroLaw). *Stern, Richard H., M-M Apr 87* 81–83
protecting hardware by copyright; printed circuit boards and their audiovisual works (MicroLaw). *Stern, Richard H., M-M Aug 87* 81, 84
protecting software simulation models of ICs from unauthorized use (MicroLaw). *Stern, Richard, M-M Oct 87* 85–89

D

Data buses
improved control acquisition scheme for IEEE 896 Futurebus draft specification. *Taub, D. Matthew, M-M Jun 87* 52–62
Data communication; cf. Local area networks
Design automation
book review; Microprocessor-Based Design, A Comprehensive Guide to Effective Hardware Design (Slater, M.; 1987). *Mateosian, Richard, M-M Aug 87* 85–86
Digital arithmetic; cf. Arithmetic
Digital integrated circuits; cf. Very large-scale integration
Digital system fault tolerance; cf. Computer fault tolerance
Digital systems
designing digital systems with LSI and VLSI circuits using SSI and MSI circuits as building blocks. *Peels, Arno J. H. M., M-M Apr 87* 66–80
Disk recording; cf. CD-ROMs
Distributed computing
Heidelberg Polyp system, reconfigurable multiprocessor for high-energy-physics experiments. *Maenner, Reinhard, + , M-M Feb 87* 5–13
high-speed distributed microcomputer system for real-time applications. *Fathi, Eli T., + , M-M Dec 87* 21-28
HM-Nucleus, distributed-kernel operating system for Homogeneous Multiprocessor, MIMD machine with shared memory facilities. *Li, Kin Fun, + , M-M Feb 87* 14–24
performance analysis methodology for Unix-based distributed file systems. *Melamed, Anna S., M-M Feb 87* 25–38
Distributed computing; cf. Local area networks

F

Fast Fourier transforms
minicomputer-based and microcomputer-based FFT implementations for advanced measurement systems. *Van der Auweraer, H., + , M-M Feb 87* 39-49
Fault tolerance; cf. Computer fault tolerance
File systems
performance analysis methodology for Unix-based distributed file systems. *Melamed, Anna S., M-M Feb 87* 25–38
Firmware; cf. Microprogramming
Floating-point arithmetic
IMS T800 transputer scientific computer with improved floating-point and communication performance; design and architecture. *Homewood, Mark, + , M-M Oct 87* 10–26

G

GaAs; cf. Gallium materials/devices
Gallium materials/devices
emulating complex instruction set computer with reduced instruction set computer using GaAs circuits. *McNeley, Kevin J., + , M-M Feb 87* 60–71
32-bit 200-MHz gallium arsenide RISC microprocessor for high-throughput signal processing environments. *Naused, Barbara A., + , M-M Dec 87* 8-20
Government – industry cooperation
European cooperation in information technology industry; Esprit program. *McLauchlan, Derek J., M-M Oct 87* 5-9

I

IEEE standards
brief update of micro standards activity (MicroStandards). *Smolin, Michael, M-M Apr 87* 92
comments on MicroStandards column in Dec. 1986 issue. *Buckley, Fletcher J., M-M Feb 87* 2
future micro standards projects (MicroStandards). *Smolin, Michael, M-M Dec 87* 88-89
generating standards in IEEE (MicroStandards). *Smolin, Michael, M-M Aug 87* 82–84
handling of draft standards in MicroStandards column. *Hill, Gary A., M-M Feb 87* 4
IEEE standards generation activities in computer area; types of standards and tabulations on Working Groups and their projects (MicroStandards). *Smolin, Michael, M-M Oct 87* 92
improved control acquisition scheme for IEEE 896 Futurebus draft specification. *Taub, D. Matthew, M-M Jun 87* 52–62
list of newly approved standards and project authorization requests in computer area (MicroStandards). *Smolin, Michael, M-M Jun 87* 85
newly approved projects in the computer area (MicroStandards). *Smolin, Michael, M-M Feb 87* 76-77
rebuttal to Letter to the Editor by F. J. Buckley concerning MicroStandards column in Dec. 1986 issue. *Smolin, Michael, M-M Apr 87* 90–91
Integrated-circuit design
protecting software simulation models of ICs from unauthorized use (MicroLaw). *Stern, Richard, M-M Oct 87* 85–89
Integrated circuits; cf. Application-specific integrated circuits

J

Japan
Japan's TRON project (special issue). *M-M Apr 87* 3–80

L

LAN; cf. Local area networks
Languages; cf. Computer languages
Large-scale integration
designing digital systems with LSI and VLSI circuits using SSI and MSI circuits as building blocks. *Peels, Arno J. H. M., M-M Apr 87* 66–80
Legal factors
analogical and 'mythological' aspects of legal reasoning (MicroLaw). *Stern, Richard H., M-M Feb 87* 73–75
Legal factors; cf. Copyright protection; Product liability
Local area networks
low-cost nonstandard LAN for controlling family of small telephone exchanges. *Micheletti, Giancarlo, + , M-M Oct 87* 70–82
Logic design
logic elements and symbol creation for designing with logic cell arrays. *Landry, Steve L., M-M Feb 87* 51–59

M

Manufacturing; cf. Product liability
Materials processing; cf. Process control
Measurement
minicomputer-based and microcomputer-based FFT implementations for advanced measurement systems. *Van der Auweraer, H., + , M-M Feb 87* 39-49
Memory management
book review; MC68851 Paged Memory Management Unit User's Manual (Motorola, Inc.; 1986). *Mateosian, Richard, M-M Apr 87* 84
capability-based microprocessor system architecture. *Corsini, Paolo, + , M-M Jun 87* 35–51

Clipper 32-bit microprocessor; system architecture, programming model, and memory management features. *Hunter, Colin B., M-M Aug 87* 6–26
memory management unit that supports demand paging and provides address mapping with overlapping rotating entries, designed for use with Unix operating system and MC68010 CPU. *Dekker, G. J., + , M-M Jun 87* 22–34

Microcomputer interfaces
80387 floating-point coprocessor architecture, interface protocols, and applications. *Perlmutter, David, + , M-M Aug 87* 42–57

Microcomputer languages
16-bit general heap processor; bit-slice and VLSI versions. *Sanchez, Eduardo, + , M-M Dec 87* 29-40

Microcomputer networks
high-speed distributed microcomputer system for real-time applications. *Fathi, Eli T., + , M-M Dec 87* 21-28

Microcomputer performance
synthetic instruction mix for evaluating microprocessor performance. *McCallum, John C., + , M-M Jun 87* 63–80

Microcomputer software, operating systems
book review; Operating Systems Design—The XINU Approach (Comer, D.; 1984). *Mateosian, Richard, M-M Oct 87* 90–91
TRON project; architecture of TRON VLSI CPU. *Sakamura, Ken, M-M Apr 87* 17–31
TRON project; BTRON, business-oriented operating system architecture. *Sakamura, Ken, M-M Apr 87* 53–65
TRON project; CTRON kernel for network nodes consisting of many kinds of computers. *Ohkubo, Toshikazu, + , M-M Apr 87* 33–44
TRON project; ITRON, industry-oriented real-time operating system architecture. *Monden, Hiroshi, M-M Apr 87* 45–52
TRON project, Japan's family of specifications for computer architectures, operating system kernels, and VLSI chips; overview. *Sakamura, Ken, M-M Apr 87* 8–14

Microcomputers; cf. Multimicroprocessing

Microprocessors
application-specific coprocessor for high-speed cellular logic operations. *Jenkins, R. E., + , M-M Dec 87* 63-70
book review; Microprocessor-Based Design, A Comprehensive Guide to Effective Hardware Design (Slater, M.; 1987). *Mateosian, Richard, M-M Aug 87* 85–86
capability-based microprocessor systems; architecture of prototype system. *Corsini, Paolo, + , M-M Jun 87* 35–51
Clipper 32-bit microprocessor; system architecture, programming model, and memory management features. *Hunter, Colin B., M-M Aug 87* 6–26
emulating complex instruction set computer with reduced instruction set computer using GaAs circuits. *McNeley, Kevin J., + , M-M Feb 87* 60–71
Japan's TRON project (special issue). *M-M Apr 87* 3–80
new generation of microprocessors (special issue). *M-M Aug 87* 4–57
synthetic instruction mix for evaluating microprocessor performance. *McCallum, John C., + , M-M Jun 87* 63–80
system considerations in design of Am29000 microprocessor. *Johnson, Mike, M-M Aug 87* 28–41
TRON project; architecture of TRON VLSI CPU. *Sakamura, Ken, M-M Apr 87* 17–31
TRON project, Japan's family of specifications for computer architectures, operating system kernels, and VLSI chips; overview. *Sakamura, Ken, M-M Apr 87* 8–14
16-bit general heap processor; bit-slice and VLSI versions. *Sanchez, Eduardo, + , M-M Dec 87* 29-40
32-bit 200-MHz gallium arsenide RISC microprocessor for high-throughput signal processing environments. *Naused, Barbara A., + , M-M Dec 87* 8-20
80387 floating-point coprocessor architecture, interface protocols, and applications. *Perlmutter, David, + , M-M Aug 87* 42–57

Microprocessors; cf. Multimicroprocessing

Microprogramming
NEC vs. Intel; implications for microcode, instruction sets, and compatibility (MicroLaw). *Stern, Richard H., M-M Apr 87* 81–83

Multimicroprocessing
book review; The Connection Machine (Hillis, W. D.; 1985). *Mateosian, Richard, M-M Apr 87* 84–85
Heidelberg Polyp system, reconfigurable multiprocessor for high-energy-physics experiments. *Maenner, Reinhard, + , M-M Feb 87* 5–13
HM-Nucleus, distributed-kernel operating system for Homogeneous Multiprocessor, MIMD machine with shared memory facilities. *Li, Kin Fun, + , M-M Feb 87* 14–24
IMS T800 transputer scientific computer with improved floating-point and communication performance; design and architecture. *Homewood, Mark, + , M-M Oct 87* 10–26
Software design for real-time multiprocessor VMEbus systems, *Heath, Walter S., + , M-M Dec 87* 71-80
SPoC, multiprocessor-based parallel programming environment. *Sterling, Thomas L., + , M-M Dec 87* 46-62

Multiprocessing; cf. Multimicroprocessing; Parallel processing

N

Natural language systems
hardware syntactic analysis processor for natural language processing. *Sanamrad, Mohammad Ali, + , M-M Aug 87* 73–80

Numerical methods
evaluation of Cordic magnification function. *Vachss, Raymond, M-M Oct 87* 83–84

Numerical methods; cf. Arithmetic

O

Object-oriented computing
capability-based microprocessor system architecture. *Corsini, Paolo, + , M-M Jun 87* 35–51
DOOM (Decentralized Object-Oriented Machine) for executing programs in POOL (Parallel Object-Oriented Language). *Bronnenberg, Wim J. H. J., + , M-M Oct 87* 52–69

Office automation
TRON project; BTRON, business-oriented operating system architecture. *Sakamura, Ken, M-M Apr 87* 53–65

Operating systems; cf. Software, operating systems

Optical recording; cf. CD-ROMs

P

Paged memories
book review; MC68851 Paged Memory Management Unit User's Manual (Motorola, Inc.; 1986). *Mateosian, Richard, M-M Apr 87* 84
memory management unit that supports demand paging and provides address mapping with overlapping rotating entries, designed for use with Unix operating system and MC68010 CPU. *Dekker, G. J., + , M-M Jun 87* 22–34

Parallel processing
DOOM (Decentralized Object-Oriented Machine) for executing programs in POOL (Parallel Object-Oriented Language). *Bronnenberg, Wim J. H. J., + , M-M Oct 87* 52–69
SPoC, multiprocessor-based parallel programming environment. *Sterling, Thomas L., + , M-M Dec 87* 46-62

Parallel processing; cf. Multimicroprocessing

Printed circuits
protecting hardware by copyright; printed circuit boards and their audiovisual works (MicroLaw). *Stern, Richard H., M-M Aug 87* 81, 84

Process control
fault tolerance and reliability for process control computers; overview and examples of European computers. *Kirrmann, Hubert D., M-M Oct 87* 27–50
TRON project; ITRON, industry-oriented real-time operating system architecture. *Monden, Hiroshi, M-M Apr 87* 45–52

Product liability
effectiveness of manufacturers' disclaimers of liability (MicroLaw). *Stern, Richard H., M-M Dec 87* 86-87

+ Check author entry for coauthors

† Check author entry for subsequent corrections/comments

Annual Index

Protocols
80387 floating-point coprocessor architecture, interface protocols, and applications. *Perlmutter, David, + , M-M Aug 87* 42–57
Publishing
book review; The Art of Desktop Publishing, 2nd edn. (Bove, T.). *Mateosian, Richard, M-M Jun 87* 84
Publishing; cf. CD-ROMs; Copyright protection

R

RD&E
European cooperation in information technology industry; Esprit program. *McLauchlan, Derek J., M-M Oct 87* 6-9
Reliability
fault tolerance and reliability for process control computers; overview and examples of European computers. *Kirrmann, Hubert D., M-M Oct 87* 27–50

S

Scientific computing
Heidelberg Polyp system, reconfigurable multiprocessor for high-energy-physics experiments. *Maenner, Reinhard, + , M-M Feb 87* 5–13
Signal processing
32-bit 200-MHz gallium arsenide RISC microprocessor for high-throughput signal processing environments. *Naused, Barbara A., + , M-M Dec 87* 8-20
Simulation
protecting software simulation models of ICs from unauthorized use (MicroLaw). *Stern, Richard, M-M Oct 87* 85–89
system performance modeling for complex VLSI; PAWS simulation language. *Iacobovici, Sorin, + , M-M Aug 87* 59–72
Software design/development; cf. Computer languages; Software development environments
Software design for real-time multiprocessor VMEbus systems, *Heath, Walter S., + , M-M Dec 87* 71-80
Software development environments
SPoC, multiprocessor-based parallel programming environment. *Sterling, Thomas L., + , M-M Dec 87* 46-62
Software, operating systems
book review; Operating Systems Design—The XINU Approach (Comer, D.; 1984). *Mateosian, Richard, M-M Oct 87* 90–91
book review; The Design of the UNIX Operating System (Bach, M. J.; 1987). *Mateosian, Richard, M-M Oct 87* 91
HM-Nucleus, distributed-kernel operating system for Homogeneous Multiprocessor, MIMD machine with shared memory facilities. *Li, Kin Fun, + , M-M Feb 87* 14–24
Japan's TRON project (special issue). *M-M Apr 87* 3–80
Software, operating systems; cf. Microcomputer software, operating systems
Software protection; cf. Copyright protection
Special issues/sections
European approaches for advanced architectures. *M-M Oct 87* 4–82
Japan's TRON project (special issue). *M-M Apr 87* 3–80
new generation of microprocessors. *M-M Aug 87* 4–57
Speech analysis/synthesis
IBM PC-XT-based speech analysis/synthesis system. *El-Imam, Yousif A., M-M Jun 87* 4–21
Standards; cf. IEEE standards
Supercomputers; cf. Multimicroprocesing
Switching systems; cf. Communication switching
Symbols
logic elements and symbol creation for designing with logic cell arrays. *Landry, Steve L., M-M Feb 87* 51–59

V

Very large-scale integration
designing digital systems with LSI and VLSI circuits using SSI and MSI circuits as building blocks. *Peels, Arno J. H. M., M-M Apr 87* 66–80
system performance modeling for complex VLSI; PAWS simulation language. *Iacobovici, Sorin, + , M-M Aug 87* 59–72
Virtual memories; cf. Paged memories

AUTHOR INDEX

A

Atwood, J. William, *see* Li, Kin Fun, *M-M Feb 87* 14–24

B

Bartels, Peter H., *see* Maenner, Reinhard, *M-M Feb 87* 5–13
Bosse, E., *see* Fathi, Eli T., *M-M Dec 87* 21-28
Bronnenberg, Wim J. H. J., Loek Nijman, Eddy A. M. Odijk, and Rob A. H. van Twist. DOOM: A decentralized object-oriented machine; *M-M Oct 87* 52–69
Buckley, Fletcher J. Standards (Ltr.); *M-M Feb 87* 2

C

Caseault, J., *see* Fathi, Eli T., *M-M Dec 87* 21-28
Chan, Ellery Y., *see* Sterling, Thomas L., *M-M Dec 87* 46-62
Chua, Tat-Seng, *see* McCallum, John C., *M-M Jun 87* 63–80
Corsini, Paolo, and Lanfranco Lopriore. The architecture of a capability-based microprocessor system; *M-M Jun 87* 35–51

D

Dekker, G. J., and A. J. van de Goor. AMORE—Address mapping with overlapped rotating entries; *M-M Jun 87* 22–34
Del Corso, Dante, and Karl E. Grosspietsch, *Guest Eds.*. Introduction to special issue on European approaches for advanced architectures; *M-M Oct 87* 4–5
Dimopoulos, Nikitas J., *see* Li, Kin Fun, *M-M Feb 87* 14–24

E

El-Imam, Yousif A. A personal computer-based speech analysis and synthesis system; *M-M Jun 87* 4–21

F

Fathi, Eli T., E. Bosse, and J. Caseault. A distributed system for real-time applications; *M-M Dec 87* 21-28

G

Gilbert, Barry K., *see* Naused, Barbara A., *M-M Dec 87* 8-20
Grosspietsch, Karl E., *Guest Ed.*, *see* Del Corso, Dante, *Guest Ed.*, *M-M Oct 87* 4–5

H

Hannum, David L. Micro Review—Graphics packages for the PC, *M-M Feb 87* 78-79
Heath, Walter, S., Software design for real-time multiprocessor VMEbus systems; *M-M Dec 87* 71-80
Hill, Gary A. Change the approach of MicroStandards? (Ltr.); *M-M Feb 87* 4
Homewood, Mark, David May, David Shepherd, and Roger Shepherd. The IMS T800 transputer; *M-M Oct 87* 10–26
Hunter, Colin B. Introduction to the Clipper architecture; *M-M Aug 87* 6–26

I

Iacobovici, Sorin, and ChakChung Ng. VLSI and system performance modeling; *M-M Aug 87* 59–72
Iseli, Christian, *see* Sanchez, Eduardo, *M-M Dec 87* 29-40

J

Jenkins, R. E., and D. G. Lee. An application specific coprocessor for high speed cellular logic operations; *M-M Dec 87* 63-70
Johnson, Mike. System considerations in the design of the Am29000; *M-M Aug 87* 28–41

K

Kirrmann, Hubert D. Fault tolerance in process control: An overview and examples of European products; *M-M Oct 87* 27-50
Kogiku, Ichizo, *see* Ohkubo, Toshikazu, *M-M Apr 87* 33–44

L

Landry, Steve L. 'Designer' logic and symbols with logic cell arrays; *M-M Feb 87* 51–59
Lee, D. G., *see* Jenkins, R. E., *M-M Dec 87* 63-70
Li, Kin Fun, Nikitas J. Dimopoulos, and J. William Atwood. The HM-Nucleus: Distributed kernel design for the Homogeneous Multiprocessor; *M-M Feb 87* 14–24
Lopriore, Lanfranco, *see* Corsini, Paolo, *M-M Jun 87* 35–51

M

Maenner, Reinhard, A fast integer binary logarithm of large arguments; *M-M Dec 87* 41-45
Maenner, Reinhard, Richard L. Shoemaker, and Peter H. Bartels. The Heidelberg Polyp system; *M-M Feb 87* 5–13
Majithia, Kenneth, ***Guest Ed.***. Introduction to special issue on the new generation of microprocessors; *M-M Aug 87* 4–5
Mateosian, Richard. Review of 'The Connection Machine' (Hillis, W. D.; 1985); *M-M Apr 87* 84–85
Mateosian, Richard. Review of 'MC68851 Paged Memory Management Unit User's Manual' (Motorola, Inc.; 1986); *M-M Apr 87* 84
Mateosian, Richard. Review of 'CD ROM 2: Optical Publishing' (Ropiequet, S., et al., *Eds.*; 1987); *M-M Jun 87* 83–84
Mateosian, Richard. Review of 'The Art of Desktop Publishing, 2nd edn.' (Bove, T.); *M-M Jun 87* 84
Mateosian, Richard. Review of 'Microprocessor-Based Design, A Comprehensive Guide to Effective Hardware Design' (Sláter, M.; 1987); *M-M Aug 87* 85–86
Mateosian, Richard. Review of 'Operating Systems Design—The XINU Approach' (Comer, D.; 1984); *M-M Oct 87* 90–91
Mateosian, Richard. Review of 'The Design of the UNIX Operating System' (Bach, M. J.; 1987); *M-M Oct 87* 91
Matsumoto, Haruya, *see* Sanamrad, Mohammad Ali, *M-M Aug 87* 73–80
May, David, *see* Homewood, Mark, *M-M Oct 87* 10–26
McCallum, John C., and Tat-Seng Chua. A synthetic instruction mix for evaluating microprocessor performance; *M-M Jun 87* 63–80
McLauchlan, Derek J. European cooperation in the information technology industry; *M-M Oct 87* 6–9
McNeley, Kevin J., and Veljko M. Milutinovic. Emulating a complex instruction set computer with a reduced instruction set computer; *M-M Feb 87* 60–71
Melamed, Anna S. Performance analysis of Unix-based network file systems; *M-M Feb 87* 25–38
Menu, Jacques, *see* Sanchez, Eduardo, *M-M Dec 87* 29-40
Micheletti, Giancarlo, and Claudio Salati. A low-cost distributed architecture for telecommunication systems; *M-M Oct 87* 70–82
Milutinovic, Veljko M., *see* McNeley, Kevin J., *M-M Feb 87* 60–71
Monden, Hiroshi. Introduction to ITRON—The industry-oriented operating system; *M-M Apr 87* 45–52
Musciano, Albert J., *see* Sterling, Thomas L., *M-M Dec 87* 46-62

N

Naused, Barbara A., and Barry K. Gilbert. 32-bit 200-MHZ gallium arsenide RISC microprocessor for high throughput signal processing environments; *M-M Dec 87* 8-20
Ng, ChakChung, *see* Iacobovici, Sorin, *M-M Aug 87* 59–72
Nijman, Loek, *see* Bronnenberg, Wim J. H. J., *M-M Oct 87* 52–69

O

Odijk, Eddy A. M., *see* Bronnenberg, Wim J. H. J., *M-M Oct 87* 52–69
Ohkubo, Toshikazu, Tetsuo Wasano, and Ichizo Kogiku. Configuration of the CTRON kernel; *M-M Apr 87* 33–44

P

Peels, Arno J. H. M. Designing digital systems—SSI and MSI vs. LSI and VLSI; *M-M Apr 87* 66–80
Perlmutter, David, and Alan Kin-Wah Yuen. The 80387 and its applications; *M-M Aug 87* 42–57

S

Sakamura, Ken, ***Guest Ed.***. Looking into the future with TRON (Special issue intro.); *M-M Apr 87* 4–6
Sakamura, Ken. The TRON project; *M-M Apr 87* 8–14
Sakamura, Ken. Architecture of the TRON VLSI CPU; *M-M Apr 87* 17–31
Sakamura, Ken. BTRON—The business-oriented operating system; *M-M Apr 87* 53–65
Salati, Claudio, *see* Micheletti, Giancarlo, *M-M Oct 87* 70–82
Sanamrad, Mohammad Ali, Koichi Wada, and Haruya Matsumoto. A hardware syntactic analysis processor; *M-M Aug 87* 73–80
Sanchez, Eduardo, Patrick Sommer, Jacques Menu, and Christian Iseli. A general heap processor; *M-M Dec 87* 29-40
Shepherd, David, *see* Homewood, Mark, *M-M Oct 87* 10–26
Shepherd, Roger, *see* Homewood, Mark, *M-M Oct 87* 10–26
Shoemaker, Richard L., *see* Maenner, Reinhard, *M-M Feb 87* 5–13
Smolin, Michael. Rebuttal: Them windmills got teeth (Ltr.); *M-M Apr 87* 90–91
Smolin, Michael. MicroStandards; *M-M Jun 87* 85
Smolin, Michael. MicroStandards—Generating standards in the IEEE; *M-M Aug 87* 82–84
Smolin, Michael. MicroStandards—More on IEEE standards generation; *M-M Oct 87* 92
Snoeys, R., *see* Van der Auweraer, H., *M-M Feb 87* 39–49
Sommer, Patrick, *see* Sanchez, Eduardo, *M-M Dec 87* 29-40
Sterling, Thomas L., Albert J. Musciano, Ellery Y. Chan, and Douglas A. Thomae. SPoC: An effective implementation of a parallel language on a multiprocessor; *M-M Dec 87* 46-62
Stern, Richard H. MicroLaw—Legal mythology, or micromyths; *M-M Feb 87* 73–75
Stern, Richard H. MicroLaw—Microcode revisited; *M-M Apr 87* 81–83
Stern, Richard H. MicroLaw—Software copyright developments; *M-M Jun 87* 81–82
Stern, Richard H. MicroLaw—Protecting hardware by copyright: Printed circuit boards and their audiovisual works; *M-M Aug 87* 81, 84
Stern, Richard H.. MicroLaw—Software models of hardware; *M-M Oct 87* 85–89
Stern, Richard H. MicroLaw—Manufacturers' disclaimers of liability; *M-M Dec 87* 86-87

T

Taub, D. Matthew. Improved control acquisition scheme for the IEEE 896 Futurebus; *M-M Jun 87* 52–62
Thomae, Douglas A., *see* Sterling, Thomas L., *M-M Dec 87* 46-62

V

Vachss, Raymond. The Cordic magnification function; *M-M Oct 87* 83–84
van de Goor, A. J., *see* Dekker, G. J., *M-M Jun 87* 22–34
Van der Auweraer, H., and R. Snoeys. FFT implementation alternatives in advanced measurement systems; *M-M Feb 87* 39–49
van Twist, Rob A. H., *see* Bronnenberg, Wim J. H. J., *M-M Oct 87* 52–69

W

Wada, Koichi, *see* Sanamrad, Mohammad Ali, *M-M Aug 87* 73–80
Wasano, Tetsuo, *see* Ohkubo, Toshikazu, *M-M Apr 87* 33–44

Y

Yuen, Alan Kin-Wah, *see* Perlmutter, David, *M-M Aug 87* 42–57

+ Check author entry for coauthors

† Check author entry for subsequent corrections/comments

MicroLaw

Richard H. Stern
Law Offices of Richard H. Stern
2101 L Street NW, Suite 800
Washington, DC 20037

Manufacturers' disclaimers of liability

A major US chip manufacturer put these notices on some of its application notes:

Life support applications

Alpha [a fictitious name] products are not designed for use in life support appliances, devices, or systems where malfunction of an Alpha product can reasonably be expected to result in a personal injury. Alpha's customers using or selling Alpha products for use in such applications do so at their own risk and agree to fully indemnify Alpha for any damages resulting in [sic, from?] such improper use or sale.

Life support policy

Alpha products are not for use as critical components in life support devices or systems without express written approval of an officer of Alpha Corp. As used herein:

1. Life support devices or systems are devices or systems which (a) are intended for surgical implant into the body or (b) support or sustain life and whose failure to perform, when properly used in accordance with instructions for use provided in the labeling, can reasonably be expected to result in a significant injury to the user.

2. A critical component is any component of a life support device or system whose failure to perform can reasonably be expected to cause the failure of the life support device or system, or to affect its safety or effectiveness [sic, adversely?].

What is this all about? Does it work?

Probably, chip manufacturer Alpha is concerned that systems designer Beta will design pacemakers, drug pumps, and the like (maybe body function monitors) that use operational amplifiers, microprocessors, or gate arrays; that Beta's system will then fail in use, leaving the user dead; and that the user's heirs and/or Beta will sue Alpha for damages. At least, the second notice addresses that possibility. The first notice might be intended to cover a little more ground. Since it does not define "life support systems," Alpha might claim that airplane or automobile navigation, brake, fuel, or anticollision devices are included in the disclaimer. However, the courts usually interpret any ambiguous language in such notices against the interests of the manufacturer who wrote the notice.

Between totally innocent John and unknowing Alpha, who should bear the cost of the loss?

Alpha has a legitimate problem, but the notice probably will not offer much protection. The problem is that Alpha does not have the least idea what systems houses will do with its chips. Alpha sets a price for its chips on the assumption that it is marketing a commodity that will go into ordinary electronic devices whose failure, while definitely a nuisance, is not fatal. Alpha is not in the insurance business, whether to guarantee the business success of its customers or that of their mutual end users. Even if it somehow were willing to go into the insurance business, Alpha would not know how to set appropriate premiums.

If Alpha tried to raise its chip prices to establish a contingency fund for lawsuits, it would have two major problems: (a) Alpha would not know how much of a fund to set up (and thus how much to raise its chip prices); and (b) if Alpha raised its chip prices, rival chip manufacturers Gamma, Delta, and Epsilon would mop up the floor with Alpha. Incidentally, some of Alpha's rivals are either offshore, judgment-proof, or unsuable. The free market will not sort things out neatly for Alpha.

On the other hand, John Q. Public is out there innocently buying the Beta pacemaker that fails, or buying a ticket on Flight 006 that strays into Russian

airspace and is shot down. In either case, the system failure may be attributed to a badly designed Alpha operator amplifier. John's heirs point out that it's not his fault that Alpha made bad op amps, adding that John had received no notice from Alpha that its defective products could risk his life. Between totally innocent John and unknowing Alpha (who does not specifically foresee the exact harm to John when an imperfect op amp fails), who should bear the cost of the loss? Especially when Alpha has the deeper pockets? In fact, John's heirs might argue that the "life support" notice is evidence that Alpha knew its product might kill John; yet, Alpha went ahead and sold the product anyway. Readers will figure out for themselves how courts and juries are likely to work that one out.

Perhaps the systems house, Beta Corp., is in a different position. It was actually warned, if it ever saw the application note.[1] Because Alpha gave notice to Beta that it specifically disclaims liability for life support systems, a court would probably say that Alpha is not liable for consequential damages resulting from chip malfunction.[2] In contracts between businessmen—rather than between a businessman and a consumer—courts usually let the parties allocate the economic risks of a sale by bargaining, at least in the absence of some type of oppression. Moreover, in the sale of goods between businessmen, the assumption is that the seller will not be liable for consequential damages resulting from risks of a type or magnitude that the seller cannot foresee. By using the disclaimer, Alpha is trying to avoid a factual controversy with Beta over what Alpha should have foreseen.

What about the provision that customers like Beta "agree to fully indemnify Alpha for any damages" from sale or use? In theory, this means that Beta agrees to reimburse Alpha for any damages that John Q. Public or his heirs collect from Alpha for pacemaker failure. Probably, such a disclaimer is no more than a lawyer's theory, or whistling in the dark. A notice on an application note is not a contract, even if Beta admits to reading it. In some states, a manufacturer might be able to get away with such a disclaimer by putting it on the back of the *invoice* for the chips, but a disclaimer on an application note is probably ineffective in any state.

In contracts between businessmen, the parties usually allocate the economic risks of a sale by bargaining. The seller is not considered liable for unforeseeable damages.

All of the foregoing discussion has been directed to a chip manufacturer's notice. Readers no doubt have seen similar notices, and more sweeping ones, on documentation accompanying software.[3] Subject to limited qualification, the same principles apply to software and hardware alike. Software marketers may argue that principles like implied warranties of fitness that apply to goods do not apply to software, which is a "service" or an "intangible," but courts do not put too much stock in that argument anymore. Whether the software or the chip sends the pacemaker haywire or dispatches the airplane to Siberia, the legal problem will probably be the same.

References

1. Beta's personnel will probably deny seeing the notice. It might have been a better idea to put the notice on all of Alpha's invoices, which Beta would find harder to deny having seen (unless Beta buys indirectly from a distributor or a jobber). Certainly, Alpha cannot fit either of those notices on the chip itself.

2. A more general disclaimer might be ineffective in some states, where courts would take the position that the chip should possess good merchantable quality and be reasonably fit for its intended purpose (whatever that is). Other states might even allow a general disclaimer of any responsibility for bad chips beyond replacement cost of the chips, but again perhaps only between businessmen.

3. Provisions of this kind are typical in "shrink-wrap licenses," which are printed provisions accompanying disks, wrapped inside shrink-wrap plastic. Typically, they state that tearing open the plastic to get at the disk constitutes user "agreement" to the terms of the license.

A recent decision in Louisiana held a shrink-wrap licensing clause unenforceable. The State of Louisiana passed a law making shrink-wrap licenses enforceable. Vault (the proprietor of the Prolok copy-protection system) sued Quaid (the proprietor of the CopyWrite program) for defeating copy protection. Quaid allegedly aided consumers to unprotect third parties' programs protected with Prolok, in violation of the third parties' shrink-wrap licenses forbidding consumers to copy the programs. Quaid also disassembled Vault's Prolok to analyze it and by reverse engineering developed a part of CopyWrite to overcome Prolok. In so doing, Quaid allegedly violated a provision in Vault's Prolok shrink-wrap license that forbade disassembly. The Louisiana shrink-wrap licensing law expressly authorized shrink-wrap licensing that prohibited "translating, reverse engineering, decompiling, [and] disassembling" computer programs.

The court held that the allegedly illegal copies of third-party software were none of Vault's business: Any complaints about violation of rights in such programs should be asserted by the proprietors of the programs, not Vault. The court then turned to the Prolok shrink-wrap license. First, the contract would be unenforceable as a "contract of adhesion" (overbearing contract), were it not for the Louisiana law favoring it. It was therefore necessary to decide whether the Louisiana law was valid. The court then held that the Louisiana law was not valid because it interfered with the operation of the federal copyright law. Accordingly, Vault was free to disassemble Prolok without liability despite the prohibitions in the shrink-wrap license.

Reader Interest Survey

Indicate your interest in this department by circling the appropriate number on the Reader Interest Card.

Low 171 Medium 172 High 173

MicroStandards

Michael Smolin
Smolin & Associates
3428 Greer Road
Palo Alto, CA 94303

Future micro standards projects

As the needs of the professional community change, so too is there a changing need for standards. When old technology and its standards lapse, new needs become paramount. On this topic, MicroStandards presents some of the ideas being discussed in the Microprocessor Standards Committee (MSC) for new standards projects.

What's in the works

Last issue's MicroStandards contained a list of the 25 currently approved standards development projects under the Computer Society's Technical Committee for Microprocessors and Microcomputers. Also listed were the 14 IEEE standards adopted from drafts prepared under the TCMM. These standards span the range from backplanes to mnemonics, from floating-point arithmetic to operating system interfaces to information transfer. The existing projects also cover BIOS, communication buses, computer languages, an instrumentation bus, and a ruggedized backplane bus. Some of the projects were authorized this year; others date back five years. The TCMM makes a continuous effort to stay abreast of the need for standards in its areas of expertise.

What's coming up, near-term

The MSC, which oversees the TCMM's standards development projects, has established several study groups. These study groups are formed when the MSC or its sponsor, the TCMM, desires to determine the feasibility of developing a particular standard before committing to a standards development project. This testing of the waters may reveal that little support exists for such a standard, that the proposal would include technology beyond the current state of the art, or that some other impediment precludes its feasibility. Conversely, the study group may attract experts in the field to help refine the scope and parameters for a standard and then work toward developing a draft. Currently, the MSC oversees four study groups:

- the Architectural Study Group (higher levels of computer architecture),
- the Superbus Study Group (performance 10 times that of Futurebus),
- the CMOS Backplane Bus Study Group, and
- the Vehicle Bus Study Group (bus for vehicle peripherals, instrumentation, controls, and sensors).

These projects are likely to be approved in the near future and represent the next range of topics subject to standardization through TCMM/MSC efforts.

MSC plans to start another project soon for the standardization of slot-function IDs (proposed as a bus-independent extension to the related Macintosh II protocol). This protocol allows for switchless automatic system configuration.

Of course, as with other standards development projects, parties outside of the TCMM and outside of the IEEE will be solicited for their comments and participation. Articles and announcements usually appear in *IEEE Micro* and *Computer* magazines and elsewhere. A broad base of interested parties participating in and contributing to the development of consensus standards has always been the IEEE ideal. Readers who are interested in participating in new projects may write directly to me. They will then be contacted when their project of interest is initiated.

What's downstream?

At this point I appeal to *you*. Where do your interests and needs lie in the area of micro standards? In what subjects would you like standards developed to aid your work efforts? (I already know about faster buses; what else?) Remember that standards can be either reactive (making order from the chaos of numerous incompatible existing approaches) or proactive (anticipating a need before the world develops numerous incompatible solutions). Guides and recommended practices are also grist for the IEEE standards mill as they also come under the aegis of the IEEE Standards Board. (See October 1987 MicroStandards for discussion of the types of IEEE standards documents.)

Please send me your suggestions. (You may volunteer to help develop the draft, but it isn't required.) I even look forward to arguments for the position that certain areas might be ill-served by standards.

Of the 33 technical committees of the Computer Society of the IEEE, only 11 are involved in standards development. The TCMM sponsors about 35 percent of the standards development projects. I will forward suggestions outside of the TCMM's areas of expertise or interest to the pertinent technical committees.

Flash!

By the time you read this, the following sponsor ballots should be under way:

- P959, I/O Extension Bus (SBX),
- P970, Advanced Backplane Bus (Versabus),
- P1096, Multiplexed High-Performance Bus Structure (VSB), and
- Reaffirmation of *IEEE Standard 796, IEEE Standard for Microcomputer System Bus* (S-100).

If you are an interested party not on the sponsor balloting body (have not received a ballot) and would like to contribute comments, contact Louise Germani at the IEEE Standards Office, (212) 705-7960.

Flash, flash!!

The MSC has asked the TCMM to cancel its project P856, Methods for Evaluating Microprocessor Performance, for lack of progress over several years. If you are interested in the maintenance of this project, make yourself known to me or Louise Germani at the IEEE Standards Office. The MSC would like to continue the project if there are any interested volunteers.

Reader Interest Survey

Indicate your interest in this department by circling the appropriate number on the Reader Interest Card.

Low 174 Medium 175 High 176

MicroReview

Continued from p. 5

cation of 417 percent is used to get a good idea of how the document will look on a 300-dot-per-inch laser printer, since the bitmapped graphics of the Macintosh screen use a density of 72 dots per inch (300/72 = 4.17). Furthermore, whatever the displayed magnification, you can select a mode in which the cursor is replaced by a magnifying glass icon. As this icon is moved over any part of the display, a small area around the cursor position is shown magnified by a factor of approximately 5.

Knuth intended TeX for the creation of beautiful books, especially books that contain a lot of mathematics. The 70 pages of *The TeXbook* that deal with producing mathematical formulas are the heart of TeX. The rest is an envelope of applied computer science worthy of much study and admiration. A macro facility allows mnemonically named and easy-to-use formatting instructions to be constructed from TeX primitives. A collection of these, called Plain TeX, is supplied with TeXtures and described in *The TeXbook*. Plain TeX allows you to achieve with simple formatting commands anything that you can produce with a WYSIWYG word processor like Word. TeX handles kerning, ligatures, varieties of spacing, and all of the niceties that have become standard in the production of fine books.

These typesetting subtleties are handled properly without your having to do anything. Also, TeX examines whole paragraphs to find optimal line breaks, and it examines pages to find the best page breaks, again automatically dealing with typesetting niceties like "widows" and "orphans."

A major benefit of TeX is portability. TeX implementations exist in a variety of computing environments, and TeX processes simple text files that can be produced anywhere and transported easily. For example, Unix tools can be used to generate TeX input files. These files can be processed by an implementation of TeX in the same Unix environment, or they can be sent over a serial line to a Macintosh, where they can be processed by TeXtures. This flexibility is unavailable with the "standard" Unix formatter, Troff.

As a tool for generating technical communications, TeXtures has several drawbacks. Like all formatters, it's slow, and the use of formatting commands is far less intuitive than the WYSIWYG user interface. Furthermore, the price is substantial, especially considering that TeX is free and all you're paying for is the Macintosh interface. Nonetheless, nothing else on the Macintosh will produce technical publications that look anywhere near as good. If you strive for perfection, then you'll want to consider TeXtures.

Reader Interest Survey

Indicate your interest in this department by circling the appropriate number on the Reader Interest Card.

Low 177 Medium 178 High 179

Continued from p. 7

Copyright Office hears screen display testimony

Richard H. Stern

On September 9, 1987, the US Copyright Office held a public hearing to determine whether display screens should be registered separately from associated computer program code. Richard H. Stern submitted a written statement and testified on behalf of the Computer Society, together with Stephen R. Levine, vice chair, Technical Committee on Computer Graphics, and Alan Paller, president of AUI Data Graphics, a division of Computer Associates International, Inc.

In addition, witnesses appeared on behalf of Apple Computer, Lotus Development, and several trade associations.

Society position

The Computer Society's position agrees with a recent Atlanta federal court decision in the case of *Digital Communications Assoc., Inc. v. Softklone Distrib. Corp.* For copyright infringement purposes, the court held that screen displays should be considered apart from both the code used to generate them and the rest of the code in the associated computer program.

The society disagreed with the present Copyright Office practice of (a) identifying the screen display with the code of the computer program, (b) treating them as a single unit, and (c) refusing copyright registration of the screen separately from the code.

The society's witnesses based their position on the propositions that

- many codes can generate the same display,
- slight changes in a code can generate a very different display,
- the type of creativity used to devise menus and other screens is different from that evoked to write the code, and
- copyright registration of screens as such would give stronger protection to proprietors of computer programs, thereby encouraging investment and creativity in the field.

Computer Society resolution

The Board of Governors of the Computer Society of the IEEE approved the following resolution on September 8, 1987:

> "RESOLVED, that the Board of Governors...in order to encourage and promote creativity and investment in this aspect of the computer graphics field, recommends that the Copyright Office...should permit an owner of a computer program screen display to register the screen for copyright protection apart from the registration of the underlying code in the computer program that generates the display."

The Board also approved presentation of written comments on this issue as submitted to the Copyright Office by Richard H. Stern. Stern is chair of the intellectual property subcommittee of the Computer Society's Committee on Public Policy (COPP) and a member of the Editorial Board of *IEEE Micro*.

Testimony

The witnesses asserted that Copyright Office refusal to register screens in this way would cause business uncertainty and insecurity due to the unpredictable scope of copyright in computer programs. "Hydraulic pressure for protection of screens" would lead some courts to protect them anyway, if not separately as screen displays, then as an aspect of the "look and feel" of computer programs. But in adopting look and feel protection, the courts might well go too far and hinder the creativity of other designers of screens. This problem could occur if the courts unthinkingly extended copyright protections to utilitarian aspects of user interfaces, which should remain in the public domain.

The witnesses expressed some differences in viewpoint. Levine voiced concern that some courts may not appreciate how much manufacturers base some screen displays on their analysis of human factors. He felt that failure to recognize that fact could lock up this approach in copyright claims, to the detriment of software progress.

Paller argued that protecting screen displays by copyright would lead to much more social benefit than possible harm. He stated that there are many ways to design good user interfaces and that providing means of capital formation is the most pressing concern in the software industry.

At the hearing Apple Computer demonstrated new screen displays based on referencing laser video disks. Apple illustrated how the authorship of screen displays differs from that of computer program code. This difference, they contend, is evidence for separate registration of screen displays and related code.

In contrast, Lotus Development—now in litigation with alleged Lotus 1-2-3 cloners Mosaic Software and Paperback Software—urged the Copyright Office to continue its present practice of registering only the underlying computer program. Lotus contended that the concept of computer program (and, by the same token, the scope of its copyright registration) should be understood broadly and flexibly. This approach would sweep up many diverse things besides the literal lines of code. Single registration of all aspects of a computer program would thus provide a flexible mechanism for protecting software that would not be tied to its present nature or the techniques for writing it.

The Information Industry Association also supported the present Copyright Office practice of allowing a single registration for all aspects of a computer program. However, they indicated no opposition to optional registration of screens as such.

IBM filed a statement taking a similar position.

Jack Russo, a private sector witness, urged the Copyright Office to create a new kind of registration for computer programs. He proposed a unitary, combined form of copyright protection of (a) the code, (b) the screen displays, and (c) all other aspects of a computer program except for the hardware. This protection would include software specifications and user interface.

The Copyright Office plans to publish its conclusions after considering the statements made by the witnesses.

Current literature

If you need up-to-date product data and pricing of standard stocked industrial electronic components, look into the *Electronic Component Catalog* from Mouser Electronics. This free, 176-page catalog offers 16,000 items, including resistors, transformers, semiconductors, hardware, integrated circuits, and microprocessors.

Mouser Electronics, 2401 Highway 287 North, Mansfield, TX 76063; (800) 992-9943; in Texas, (817) 483-4422.

Informer Computer Terminals, Inc., offers the *Dial-Up Networking Primer* as a conceptual guide for both technical and nontechnical readers. This illustrated primer is a quick reference or a more complete data communications guide, depending upon need. The booklet discusses company products, dial-up networks, and user implementation requirements.

Informer Computer Terminals, Inc., 12781 Pala Drive, Garden Grove, CA 92641; (714) 891-1112; free upon request.

Howard W. Sams has recently published three computer-oriented tomes—one for would-be desktop publishers, another for Unix aficianados, and yet another for beginning/intermediate Mac programmers.

Personal Publishing with PC PageMaker reveals techniques for designing publications from one-page flyers to multipage journals. It provides hands-on instruction, numerous visual examples, and an explanation of typesetting terms. Topics include selecting the right hardware and software, assembling a personal publishing system, working with type, building master pages, and adding graphic elements.

Author Terry Ulick also wrote *Personal Publishing with the Macintosh* and founded several magazines on the subject.

Unix Communications is a reference guide for users who need to communicate with other users and other systems. The book follows on standard documentation with coverage of Unix mail, networking, and file transfer tools. Authors Bart Anderson, Bryan Costales, and Harry Henderson of the the Waite Group demonstrate use, control, and programming of Unix communication tools with examples.

MPW and Assembly Language Programming for the Macintosh is an introductory-level book on the Macintosh Programmer's Workshop that teaches the Macintosh assembly language. Programmers can write programs using the Macintosh Toolbox. Mouse events, windows, QuickDraw, and menus are covered. This tutorial includes the MPW shell command language and the 68000 Instruction Set with directives and toolbox traps. Author Scott Kronick is a core member of the Berkeley Macintosh Users Group.

Howard W. Sams & Company, 4300 West 62nd Street, Indianapolis, IN 46268; (317) 298-5400; 320 pp., $18.95; 350 pp., $26.95; and 352 pp., $24.95, respectively.

MicroTidbits

Now you can record directly over magneto-optic data, thanks to a new technique developed at **Carnegie Mellon's Magnetics Technology Center.** Before this discovery, magneto-optics users had to erase before they could rerecord. Mark Kryder et al. have eliminated the erasure step by using a temperature-change method to write directly over previous data.

AT&T and **DEC** researchers are developing test methods with the **National Bureau of Standards** for the IEEE Standard for Portable Operating System Environments, commonly known as Posix. Industry and users need test methods to ensure that new operating system environments conform to the proposed Posix standard. NBS is also developing a Federal Information Processing Standard based on Posix for federal government use.

Access Data Products has set up an electronic bulletin board entirely dedicated to PC connectivity and networking. The board operates 24 hours a day at no charge with dedicated phone lines at (914) 667-1841, (914) 667-1842, and (212) 319-7300.

A monthly newsletter for emulator users has hit the mails. *News & Notes* provides educational articles, new product information, product updates, and a special emulator clinic column. To obtain a free subscription, contact Tammie Adams, **Softaid, Inc.**, 8930 Route 108, Columbia, MD 21045; (301) 964-8455 or (800) 433-8812.

Brave new PCs projected

Can you believe that in five years a typical office system will have 10 times the computing power, 12 times the main memory, and nearly 100 times the mass storage capacity of current PCs?

The editors of the *Computer Industry Almanac* predict that a system comparable to today's $3000 model will boast a 32-bit microprocessor, 8M bytes of main memory, one 3.5-inch floppy-disk drive, and one hard-disk drive with 80M bytes of storage. They also envision a high-resolution color graphics display, a 9600-baud built-in modem, and a built-in LAN interface. The system printer could produce letter-quality, high-speed draft printing of text and graphics. This dream machine's performance could be as high as 3 MIPS.

A budget-model, $1000 system would be a scaled-down version of today's office PC with CD-ROM for recreational applications. It too would be fired by a 32-bit microprocessor but would harbor only 4M bytes of main memory. One 3.5-inch floppy-disk drive, a color display, one CD-ROM drive, and a letter-quality printer round out the possibilities. Its performance is estimated at 1 to 1.5 MIPS.

In addition to product forecasts, the 780-page *Computer Industry Almanac* includes an industry overview and rankings of 400 companies, 1000 key people, and 800 top products. Current data on advertising, salaries, and sales in the computer industry offer points of comparison.

Buy the almanac in book stores or from the publisher (add $2.00 for shipping): Computer Industry Almanac, Inc., 8111 LBJ Freeway, Dallas, TX 75251-1313; (214) 231-8735; hard cover $49.95, paperback $29.95.

Reader Interest Survey

Indicate your interest in this department by circling the appropriate number on the Reader Interest Card.

Low 183 Medium 184 High 185

New Products

Marlin H. Mickle
University of Pittsburgh

Send announcements of new microcomputer and microprocessor products, and products for review, to Managing Editor, IEEE Micro, *10662 Los Vaqueros Circle, Los Alamitos, CA 90720-2578.*

Color images transmitted via telephone lines

Color Video Fax Corporation of America is marketing the Japanese Analogue & Digital Science system for transmission of still, color images. According to the company, its Eris system transmits 512 × 480-pixel × 8-bit images via standard telephone communication lines in less than one minute.

The Eris digital transmission contains automatic error correction, automatic sending and receiving capabilities, and remote control from personal computers for system expansion. Pictures can be sent over telephone lines with optional use of a television monitor, video camera, VHS recorder, laser disk, or computer. Pictures can be stored on a floppy disk and/or output in hard copy on Polaroid, Hitachi, and Sharp printers.

Contact the company for pricing.

Reader Service Number 1



Color Video Fax Corporation of America markets Japan's Eris color-image, digital transmission system.

Hypermedia organization comes to Apples

Apple Computer's HyperCard allows Macintosh users to organize, use, customize, and create new information with multiple information types such as text, graphics, video, music, voice, and animation. An English language-based scripting language allows users to write their own programs.

With HyperCard, users can organize and use information the way they think—by association, context, and hierarchy—while browsing and searching through large bodies of information. Users sort, make notes, type, or draw on these cards in the same way they might use paper index cards.

Upon activating a button, users can link a card or part of a card to another card or a stack of cards. Buttons also perform tasks such as dialing the telephone, sorting a stack, or finding a videodisc sequence.

HyperCard consists of a main disk, help disk, stack examples and ideas disk, and backup; it requires a 1M-byte RAM Macintosh with two 800k-byte floppy-disk drives.

Suggested retail price of Apple's three-disk HyperCard is $49. It will be included with all new Macintosh computers.

Reader Service Number 2

VLSI diagnostic workstation announced

The Schlumberger/ATE Integrated Diagnostic System IDS 5000 Workstation tests and diagnoses VLSI technology. The workstation combines scanning electron microscope technology with CAD/CAE tools.

With the IDS 5000, logic designers can observe behavior of a circuit while simultaneously monitoring circuit connectivity and predicted behavior through the schematic netlist, layout database, and simulation tools.

Delivery of the Schlumberger/ATE IDS 5000 is 60 days ARO. The list price is $495,000, including a one-year warranty.

Reader Service Number 3

DSP simulators run on IBM PCs

Software from Avocet Systems simulates and debugs Texas Instruments signal processing chips. AVSIM321 and 322 used on IBM PC compatibles work with the TMS 320 family of chips.

In operation the software produces registers, flags, and program and data memory on screen for manipulation by editing keys. Function keys control program flow for debugging. An Undo key uses the simulator's trace memory to back up one-at-a-time through recently executed instructions for error detection. Screen menus and a command mode are added features.

The Avocet Systems simulator/debuggers are priced at $379.

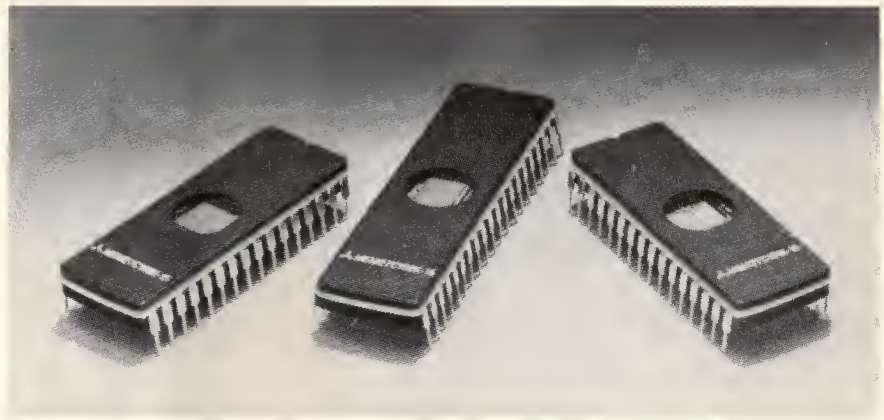
Reader Service Number 4

Mitsubishi offers three 150-ns, 1M-bit EPROMs

Mitsubishi Electronics America expanded its UV EPROM product line with three 150-ns, 1M-bit CMOS devices. Available in both 128K × 8- and 64K × 16-bit organizations, the EPROMs store nonvolatile memory in applications using 16/32-bit microprocessors.

Two of the devices, the M5M27C100K and M5M27C101K, are available in a 32-pin Cerdip package. The M5M27C-100K is pin compatible with the existing generation of 28-pin, 1M-bit mask ROMs. The M5M27C101K conforms to the JEDEC standard pinout for byte-wide memory.

A third device, the 64K × 16 M5M27-C102K, is encased in a 40-pin Cerdip package. All three offer a page programming algorithm that allows four bytes to be programmed at the same time.



Three CMOS EPROMs from Mitsubishi Electronics America.

Mitsubishi offers the devices for $40.75 each in 100-piece quantities for the 150-ns speed grade.

Reader Service Number 5

AT&T desktop computer uses 80386 chip

The 6386 WorkGroup System is part of AT&T's recently announced computer and data networking product line. The 6386 uses the Intel 80386 chip and includes a nonproprietary operating environment that lets it share software and data with the rest of the line and with other computers that use standard operating systems.

The 32-user system runs 32-bit Unix System V applications concurrently with MS-DOS applications and features a DOS Supervisor for running up to eight applications simultaneously at high speeds.

A 6386 desktop model includes 1M-byte RAM expandable to 48M bytes; a floor-standing model, the 6386E, includes 2M-byte RAM expandable to 64M bytes. Both models use either 5.25-inch floppy disks or 3.5-inch minifloppies.

List prices for the AT&T 6386 WorkGroup System begin at $4899.

Reader Service Number 6

CAE tools aid power circuit designs

Analog Design Tools has announced a set of design tools for the Analog Workbench that is focused on the simulation and analysis needs of designers working on power supplies and other power circuits.

The Power Design Tool Kit combines a nonlinear model of transformer core magnetics with libraries of magnetic core materials and semiconductor power devices. Users can select tolerance distribution functions for statistical analysis calculations. Air gap effects on core behavior can be modeled, and semiconductor thermal resistivities can be modified for stress analysis.

The Analog Design Tools software for the Analog Workbench with three libraries costs $33,000.

Reader Service Number 7

Network controls 1600 RS-232 devices

The CMCNETII communications network from Connecticut Micro-Computer enables an RS-232 computer or terminal to control a network of up to 1600 RS-232 devices at distances up to several miles. The network is recommended for factory automation and information gathering systems.

The system consists of four types of modules, a controller attached to the PC, device modules connected to each RS-232 device, expansion modules for systems of over 40 devices, and a repeater module for driving branches of the network more than 4000 feet.

CMCNETII modules cost under $100 in OEM quantities.

Reader Service Number 8

CMOS static RAMs feature 35-ns access times

Three CMOS static RAMs from Advanced Micro Devices are available in standard and low-power versions for commercial and military applications. Applications for the devices include cache and buffer memory, graphics, distributed processing, multiprocessing system, disk controllers, and instrumentation.

The 16K × 4-bit Am99C164 and Am99C165 operate at 605 mW maximum in standard-power versions and offer standby power at CMOS input levels of 82 mW. The low-power versions consume 495 mW operating power and 1.6 mW standby power at CMOS input levels. Their access time for military use is 45 ns.

The Am99C164 features chip enable and write enable controls for array application mode; the Am99C165 incorporates both functions and adds an output enable control, alleviating bus contention in high-speed systems with large memory banks.

The Am99C88H features access times of 35 ns for commercial use and 45 ns for the military version. The device is offered in standard and low-power versions. All three devices complement the Am99C641 64K × 1 static RAM.

In 100-piece quantities the DIP and leadless chip carrier devices are priced from $20 each.

Reader Service Number 9

Optical system design aid announced

Dynacomp has introduced three educational and scientific software packages for IBM PCs and compatibles. Optics One, an optical ray-tracing program, employs recursion and iteration schemes to minimize both truncation and round-off error. The menu-driven lens design software allows users to zoom in on a graphical display to fine-tune their designs. The positions of surfaces, locations and heights of images, and Petzval sums are included in the general report. Optics One requires a 256K RAM, MS-DOS 2.0 or higher, and a graphics card.

Plotsmith is another menu-driven package for data plotting up to 10 sets of 250 data points each. Users enter and edit data from the keyboard.

Eigen Analyzer, designed for Apple, CP/M, and IBM computers, helps users find the real and complex eigenvalues of a real matrix, whether specified uncertainties exist in the matrix elements or not. Users can enter, edit, and save elements for later recall from disk as ASCII files.

Dynacomp sells Optics One for $99.95 (disk plus manual) and Plotsmith and Eigen Analyzer for $49.95 each (disks).

Optics One **Reader Service Number 10**
Plotsmith **Reader Service Number 11**
Eigen Analyzer **Reader Service Number 12**

Software Engineers

Northrop Corporation's Defense Systems Division, located in sprawling Rolling Meadows, IL just northwest of Chicago, provides a state-of-the-art software development environment implemented on a VAX cluster configuration, running under VMS, connected to Sun work-stations on an Ethernet fiber optics LAN, running under UNIX. Each software engineer has a terminal with access to any system on the network. Terminals are being replaced by personal workstations.

We offer professionals with a BSCS, BSEE, BS Math or Physics (or equivalent) MS preferred, and a minimum of 3 years experience, opportunities in the following areas. **Management, Systems Architect, Technical Leaders and engineering assignments available.**

Systems Programmers

Our many, varied applications require significant growth in our support capabilities. We need the best people with experience in:

- LANGUAGES, including Ada, Assembler, C, FORTRAN, JOVIAL, and Pascal
- OPERATING SYSTEMS, including UNIX and VMS
- Development of Real-Time Operating Systems
- Development of Software Tools
- Performance Modeling and Evaluation
- Use of Structured Software Development Methodologies

Software Systems Engineers

Our software engineers develop software from systems requirements through implementation, and need experience in:

- Software Requirements Analysis
- Architectural Design
- Software Validation and Test Specification
- Performance Specification and Modeling
- Interface Design and Specification

ECM/EW Systems Software Engineers

ECM/EW Systems are our business. Experience should include:

- Real-Time Control Systems
- Radar Data Processing
- Embedded Computer Systems
- Systems and Unit Level Diagnostics
- Object Discrimination & Classification
- ECM Algorithm Development
- Kalman Filtering
- Optimal Control

Hardware Diagnostics Software Engineers

We design and develop advanced systems using the latest hardware and software technologies for our military clients. Experience required in:

- Intelligent Control Panel Systems Development
- Built-in-Test, Functional Test
- Micro and Macro Diagnostics for Fault Identification

Interested individuals are encouraged to forward resume to: **Supervisor-Staffing, Dept. C98, Northrop Corporation, Defense Systems Division, 600 Hicks Road, Rolling Meadows, IL 60008.** An equal opportunity employer M/F/V/H. U.S. Citizenship required.

NORTHROP
Defense Systems Division
Electronics Systems Group

PS/2 software links laptops, PCs

Meridian Technology's Carbon Copy Plus data communications program is available on 3.5-inch disks for use on IBM's PS/2 family and compatible portable laptop computers. The software combines PC-to-PC remote control, PC-to-host terminal emulation, and X-modem and Kermit file transfer protocols in an integrated package.

In remote-control mode, the program joins together two PCs over an asynchronous link so that a keystroke entered on one appears simultaneously on the other. In a second mode, Carbon Copy Plus emulates DEC VT-52 and VT-100, Televideo 920, and IBM 3101 asynchronous terminals. Users can access host computers and on-line information databases.

The Meridian Technology software costs $195.

Reader Service Number 13

VMEbus board supports TOP protocol

Motorola/Microcomputer Division's MVME374 controller board hosts an Ethernet chip set and supports the seven-layer Technical and Office Protocol with 32-bit data transfer.

The VMEbus-based board also contains 1M byte of 32-bit-wide shared DRAM with parity for one-wait-state access to RAM supporting the MC68020 MPU and zero-wait-state cycle access for Ethernet Lance chip use.

The accompanying MicroTOP 1.0 protocol software is based on current OSI software and patterned after the ITI-certified MicroMAP 2.1 protocol stack.

Motorola expects the MVME374 to reach production early in 1988 and sell for $1795. Software object codes should sell for $600 each and source codes for $50,000.

Board **Reader Service Number 14**
Software **Reader Service Number 15**

PS/2 tape systems store up to 600M bytes

Mountain Computer has expanded its line of tape backup and disk data storage with four products that support IBM Personal System/2 models.

Two tape systems operate under the DOS 3.3 operating system and include Advanced Version 4.4 tape backup software. Both are compatible with IBM PC Net, Token Ring, Novell, and Xenix System 5 SCO networks and store from 40M bytes to 600M bytes of data.

Two disk drive systems are available for immediate shipment. One is a 50M-byte version of DriveCard, which plugs into the PS/2 Model 30 expansion slot. The other is a dual-drive version of Mountain's MicroBernoulli. An external unit, it provides two removable, 20M-byte cartridges and operates at hard-disk speeds.

Tape unit prices begin at $595; disk drives, at $1095.

Tape units **Reader Service Number 16**
Disk drives **Reader Service Number 17**

Neurocomputer adds 80386 processing speed

The Anza AZ1500 neurocomputer system from Hecht-Nielsen Neurocomputer Corp. offers neural network researchers and application developers an 80386-based host computer. Typical preprocessing algorithms that benefit from the added speed include digitizing, normalization of data, and Fourier transforms. Applications include machine vision, speech recognition, and real-time signal analysis.

The host computer is a Zenith 386-80, which acts as an I/O device for the AZ1500 IBM and compatible neurocomputer coprocessor board. The neurocomputer implements virtual neural networks containing up to 30,000 processing elements with up to 480,000 interconnects. Included with the system is the HNC Neurosoft software of five Neural Net packages and User Interface Subroutine Library.

The Anza AZ1500, with neuro-computing coprocessor integrated into a Zenith 386-80 and a 13-inch color monitor, is available immediately for $19,500.

Reader Service Number 18

Lab, factory products offered for Macintoshes



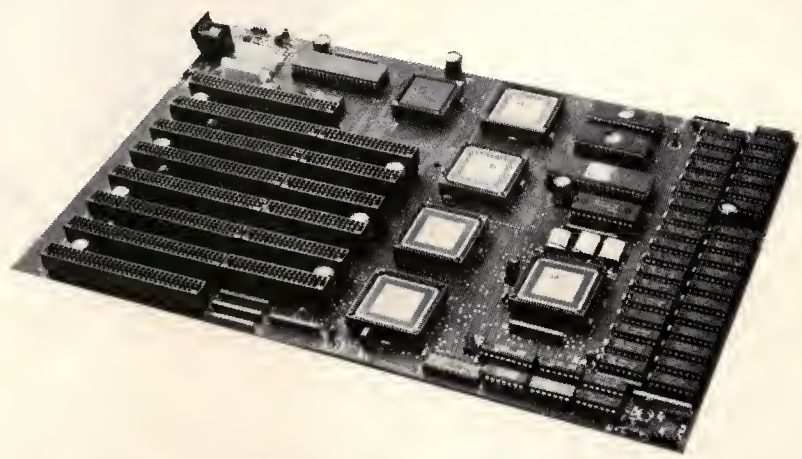
Six Macintosh II boards, two SE boards, and application software form the Analog Connection laboratory/factory series from Strawberry Tree Computers.

The Analog Connection series from Strawberry Tree Computers contains application software and eight plug-in boards to help users acquire and control data in laboratories and factories when using Apple's Macintosh II or SE computers. WorkBench application software ties together the boards, which provide 16-bit and 12-bit resolution and up to 16 analog input channels and 16 digital I/O lines.

Users select analog/digital inputs and outputs, data loggers, and chart recorders and tie them together on the screen to create symbolic representations of what actually happens in hardware. Typical use would include the measurement of voltage, temperature, pressure, strain, weight, and other parameters and control of heaters, boilers, pumps, valves, and other laboratory and industrial equipment.

Analog Connection WorkBench software costs $495; the boards begin in pricing at $595.

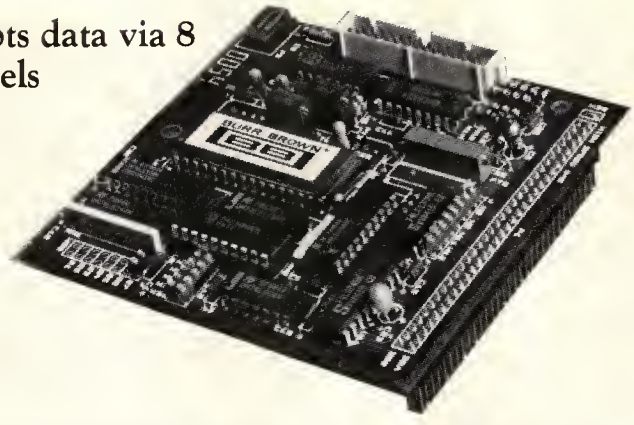
Software **Reader Service Number 19**
Boards **Reader Service Number 20**



VLSI Technology's VL82PCAT-SB1 lets users evaluate the features of the company's VL82CPCAT-QC chip set, which replaces PC/AT devices. The module contains a 12-MHz 80286 microprocessor and 36 120-ns, 256K × 1 dynamic RAMs for one-wait-state reads and writes at 12 MHz. $2000 each.

Reader Service Number 21

Module accepts data via 8 analog channels



Burr-Brown's Data Acquisition Module features a maximum multiplexer settling time of 3.5 μs, an A/D conversion time of 4 μs, an aperture jitter of 0.3 μs, and a total conversion time of 5.5 μs.

Burr-Brown has announced a 180K-channel-per-second data acquisition module designed to operate with the plug-in PCI-20000 Personal Computer Data Acquisition System.

The PCI-20023M-1 module provides eight single-ended analog-input channels and can be expanded to 128 channels. Intended for high-level signals, the module incorporates a sample/hold and A/D converter to provide high speed sampling rates. Internal hardware can configure the module to automatically increment to the next channel after each Start Convert command. Conversions can be started by the system's rate generator, an external signal, upon reading the previous conversion, or by software command.

Unit price of the PCI-20023M-1 is $895; delivery is from stock.

Reader Service Number 22

Unix coprocessor supports 32-bit computers

The Aeon:332/AT from Aeon Technologies is a coprocessor for the IBM AT and compatibles based on National Semiconductor's 32332 microprocessor. Designed for real-time and Unix applications, the board operates at 15 MHz with up to 4M bytes of no-wait-state DRAM. A 32382 demand-paged MMU allows full use of the 32-bit virtual address space.

Features include bus master capability, access to the AT-bus address space, and two 16-bit counter/timers. Operating system support includes Aeon5.3, Hunter and Ready's VRTX, and stand-alone access to DOS and BIOS functions. Languages available for Aeon5 on the 332/AT are C, Pascal, Fortran, Lisp, Basic, and Assembler.

Full-screen editors, symbolic source-level debuggers, and source code managers are included with the Aeon:332/AT.

The Aeon Technologies coprocessor is available with 1M byte of DRAM in OEM quantities 30 days ARO for $1995.

Reader Service Number 24

Digital sound converted to speech, music

The Covox Speech Thing is a digital sound converter that attaches to an IBM PC or compatible parallel printer port to produce speech, sound effects, and music for business, education, and pleasure.

Speech Thing includes an audio amplifier, built-in speaker and headphone jack, manual, and software for use as a resident talking appointment calendar, English/Spanish calculator, and game demonstration. Additional software includes a graphics-based sound editor, special-effects control panel, and prerecorded vocabularies for inclusion in user-written programs. Smooth Talker, the First Byte text-to-speech synthesizer, comes with Speech Thing; the Voice Master speech synthesizer and digitizer for voice recognition and stereo sound (used to prepare Speech Thing software) is available separately.

Covox sells Speech Thing for $69.95.

Reader Service Number 23

VME engine promises zero wait state



Heurikon Corporation's HK68/V2F features a VSB-compatible memory expansion bus, mailbox interrupt support, and one RS-232 serial port.

Heurikon Corporation has announced a 68020-based VME engine with 20-MHz, zero-wait-state operation for reads and one-wait-state for writes. Designed for real-time processing, the HK68/V2F offers up to 4M bytes of on-board DRAM with parity, 128K EPROM, 128 bytes of nonvolatile RAM, and optional 68881 floating-point coprocessor.

The HK68/V2F is available at 20-MHz operation for $1495 in OEM quantities of 100. A no-wait-state version is available for $2395 in the same quantities.

Reader Service Number 25

Handheld analyzer contains disk drive

Network Communications Corporation has developed a handheld network diagnostic instrument that includes a built-in, 3.5-inch floppy-disk drive. Its 50K to 700K data partitions, written in real time, can help users track down diagnostic problems by allowing them to compare before/after data or archive separate diagnostic tests. Users can also store programs, set up information on disk, and store or mail the minidisks.

Functions included in the 5-pound 6610D Network Probe are data line monitoring, protocol analysis, bit and block error-rate testing, RS-232C lead status analysis and visual display, and asynchronous terminal operation.

NCC sells the 6610D for $3800.

Reader Service Number 26

Atari ST scanner is menu driven

The ST Image Scanner from Navarone Industries inputs photographs, graphics, and technical drawings into the Atari ST computer via a Canon IX-12 interface.

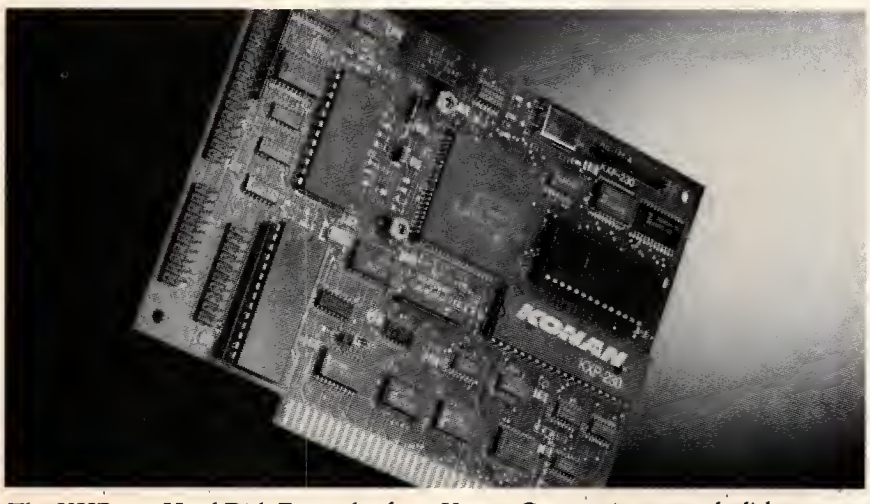
According to the manufacturer, the ST Image Scanner digitizes images in halftone mode for photographs or line-art mode for drawings and logos in less than 15 seconds. The entire image or selected area can be captured with resolutions of 75 to 300 dpi and 32 shades of gray.

The scanner software uses drop-down menu commands and is compatible with graphic programs that edit, crop, size, and print the image.

With interface, cable, and software, the Navarone Industries ST Image Scanner costs $1239.

Reader Service Number 27

PC hard disks expanded to 302M bytes



The KXP-230 Hard Disk Expander from Konan Corporation extends disk capacity and improves drive access and file loading times through data, directory, and FAT caching.

A hard-disk controller that Konan Corporation says will increase storage in IBM PC, Tandy, and compatible computers by 100 percent is now on the market.

Model KXP-230 Hard Disk Expander uses data compression and file compaction techniques to compress individual files of highly repetitive data by as much as 800 percent. Data is stored on the disk in normal Modified Frequency Modulation format, eliminating the need for special RLL-certified hard disks when the KXP-230 is used. The board supports hard disks or volumes up to 302M bytes.

KXP-230 contains a disk cache for storing often-used data, disk correction for errors up to 4096 bits in length, and fragmentation control for contiguous file storage.

Suggested retail pricing of Konan Corporation's KXP-230 is $249 for use on PC, XT, or Tandy 1000 computers and $299 for PC ATs.

Reader Service Number 28

CLASSIFIED ADS

RATES: $5.00 per line, $50 minimum charge (up to ten lines). Average six typeset words per line, nine lines per column inch. Add $4 for box number. Send copy at least one month prior to publication to: **Heidi Rex** or **Marian Tibayan, Classified Advertising, IEEE MICRO Magazine, 10661 Los Vaqueros Circle, Los Alamitos, CA 90720; (714) 821-8380.**

In order to conform to the Age Discrimination in Employment Act and to discourage age discrimination, *IEEE Micro* may reject any advertisement containing any of these phrases or similar ones: "...recent college grads...," "...1-4 years maximum experience...," "...up to 5 years experience...," or "...10 years maximum experience." *IEEE Micro* reserves the right to append to any advertisement, without specific notice to the advertiser, "Experience ranges are suggested minimum requirements, not maximums." *IEEE Micro* assumes that, since advertisers have been notified of this policy in advance, they agree that any experience requirements, whether stated as ranges or otherwise, will be construed by the reader as minimum requirements only.

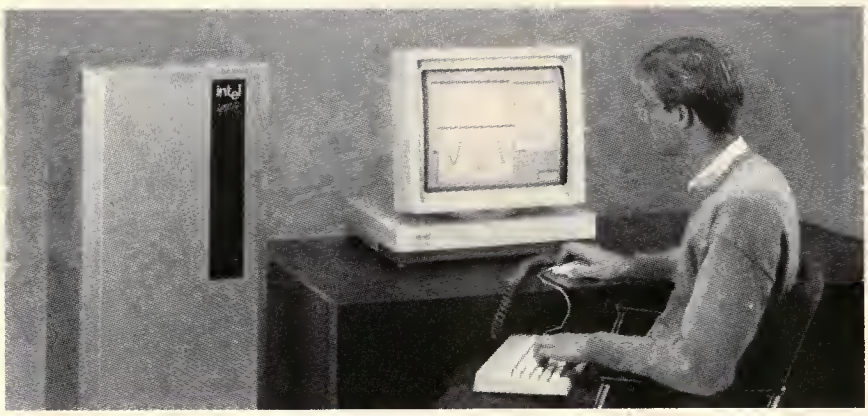
U.S. NAVAL ACADEMY
at Annapolis, Maryland

Seeking individuals qualifying for Assistant Professor position (tenure track).
Qualifications: An earned doctorate coupled with a strong background in the use of computing in the classroom.
Responsibilities: Assists faculty members with the integration of computer technology into the engineering curriculum. Must teach minimum of one undergraduate course per semester in the area of personal expertise.
Salary: Negotiable depending on experience and qualifications.
Personal expertise is expected to fall within one of the following Engineering disciplines: Aerospace, Electrical, Mechanical, Naval Architecture, Ocean, Marine, or Weapons & Systems.
Position is available June 1, 1988.
Closing date: December 31, 1987.
Submit letter of application, resume and minimum of three references to:
Director, Computer Services
Ward Hall
United States Naval Academy
Annapolis, Maryland 21402-5045
Attn: Mr. Doug Afdahl (301) 267-3693
(Equal Opportunity Employer).

WESTERN OREGON STATE COLLEGE
Computer Science Professor

Assistant professor to teach upper and lower division courses in operating systems, computer organization and data structures. Earned Doctorate in Computer Science required. A record of successful teaching experience preferred. 9-month, tenure-track appointment, $23,000 minimum salary. Send letter, curriculum vitae, and three letters of recommendation by February 1, 1988 to: Dr. David Olsen, Search Committee, WOSC, Monmouth, OR 97361, (503) 838-1220, ext. 509. Western is a liberal arts college of 3,650 students, which places special emphasis on undergraduate teaching. An Affirmative Action, Equal Opportunity Employer.

Second-generation hypercube offered



Intel's iPSC/2 system supports a fluid dynamics and heat transfer software package called Nekton from Nektonics, Inc. The numerical simulation tool solves complex 2D and 3D problems.

The Intel Scientific Computers second-generation hypercube features new hardware and software developments for easier programming and faster speeds. The iPSC/2 family of concurrent supercomputers can be used as compute servers for large-scale scientific and engineering applications.

Standard iPSC/2 systems are available in configurations from 16 to 128 processing nodes, with up to a gigabyte of memory. Vector concurrent iPSC/2 VX systems, with a vector arithmetic acceleration at each node and with up to 64 nodes, promise a peak performance of 400 MFLOPS.

Each system node features a 4-MIP 80386 and 80387. Surface-mounted 1M-bit DRAM modules provide 1 to 16M bytes of memory per node. Concurrent Workbench programming permits simultaneous access to the iPSC/2 system from a network of engineering workstations through a System Resource Manager interface, an 80386-based Unix V.3 system.

iPSC/2 pricing begins at $200,000. Original system upgrades are available.

Reader Service Number 29

Managers can forecast patent trends

Patents-PC software from Battelle allows corporate managers to track large quantities of data from the US Patent and World Patent Index databases. This data can be used to forecast technology developments, perform competitive analysis, and determine specific trends.

The software package includes a personal computer disk that can be copied onto a hard disk, sample database disk, user's guide and reference manual, and two hours of telephone consulting time with company experts. Intended for an IBM or compatible computer with a hard disk and 640K memory, the software is also compatible with Hercules, CGA, or EGA video monitors.

A site license for Patents-PC is available for $7500 for the first locale.

Reader Service Number 30

NS 32000's gain real-time operating system

Industrial Programming has announced a MTOS-UX real-time, multitasking operating system for National Semiconductor's 32000 series of microprocessors. According to the company, the MTOS-UX/32K operating system has the same internal structure and user inferface as other MTOS-UX computers, and programs written in a high-level language run in the same way.

A link to Unix permits a system or group of systems running under MTOS-UX to run as companions to a Unix-based workstation, in the manner of a work cell and its controller. It can log on and load object modules from files prepared under Unix.

Contact the company for pricing.

Reader Service Number 31

PC multiprocessor supports 64 users

Corollary Inc. has unveiled plans for a generic multiple processor system designed for 386 PC compatibles. Called the Attain 386, the system distributes user workload among up to five 386 processors, supporting up to 64 users.

Attain 386 includes an 80386 processor, 4M bytes of memory, and four specialized I/O ports that each interconnect to the company's current eight-line terminal concentrator. A 386 system can be configured with up to four Attain 386 processors, while maintaining SCO Xenix 386 binary compatibility for interchangeable, generic computing.

Corollary's plans call for December shipments; no pricing has yet been set.

Reader Service Number 32

Raster editor supports scanner

Houston Instrument's Hi-Scan Raster Graphics Toolkit enables users to edit and manipulate drawings scanned into a computer with the company's Scan-CAD plotter accessory. An icon-driven utility package, Hi-Scan works with up to E-size raster image files, allowing manual conversion to vector files or hard-copy output. Either a True Grid or Hipad digitizer, a mouse, or the keyboard arrow keys can be used to control the cursor.

Houston Instrument will mail Hi-Scan software to users who have purchased a Scan-CAD Model 128 and have applied for warranty registration.

Reader Service Number 33

Macintosh II drive accesses data in 26 ms

The newest member of the CMS Enhancements PRO-Series family of hard-disk subsystems is a 140M-byte internal drive designed for Apple Computer's Macintosh II.

Available immediately, the 5.25-inch, half-height PRO-140 II/i features an average access time of 26 ms. The subsystem includes the CMS SCSI Utilities program that assists in formatting, initializing, and installing the disk drive. Suggested retail price is $1695.

Reader Service Number 34

Digitize nondimensioned parts faster

PMX's Digi-Graph provides productivity improvements of up to 10 to 1 in digitizing complex, nondimensioned machine development parts and artwork. The software is useful for digitizing such geometries as die blank developments, artistic detailing, lettering, and logos.

With Digi-Graph users input straight-line moves, true arcs, and rapid motion statements. Lines are defined by their endpoints, and arcs by any three points on the circumference.

The IBM PC software creates DXF-formatted files for data exchange between XL/NC NC Programming software, AutoCad, Cadkey, and other DXF systems. The digitizing pads supported include GTCO, Pnumonics, Hewlett Packard, Kurta, Houston Instrument, and Skalar Systems.

Contact the company for price.

Reader Service Number 35

A/D converters offer 20-kHz sampling rates

Two 16-bit sampling A/D converters from Micro Networks can be used to digitize high-frequency signals in digital signal processing applications.

The MN6290 and MN6291 converters are constructed with internal, user-transparent, track-hold amplifiers. These amplifiers enable the units to sample and digitize dynamically changing input signals (with frequency content up to 10 kHz) at sampling rates up to 20 kHz.

MN6290 has a 10-volt input span; MN6291's span is 20 volts. Each features a 40-μ maximum conversion time, a 5-μ acquisition time, and a 10V/μ slew rate with precision of $\pm$ 0.001 percent linearity error. Both devices offer two electrical grades and two operating temperature ranges.

Pricing for each device begins at $252 for quantities up to 14; production quantities require 12 to 16 weeks for delivery.

Reader Service Number 36

Software analysis tools announced by Intel

Intel Corporation has introduced 80286 and 8086/8088 real-time software analysis tools that can be run on the IBM PC AT and XT. Called iPAT-286 and iPAT-86/88, these products allow software developers to monitor the execution of real- or protected-mode software for speed tuning and test analysis.

iPAT tools permit the host system to display high-level histograms, tables, and code coverage maps and examine the code generated by 8086/80286 compilers and assemblers. The analyzers use high-level-language symbolics (ASM, C, PL/M, Pascal, Ada, and Fortran) so designers can quantify code behavior at the module, procedure, statement, or address levels.

Intel's analyzers begin in price at $2995; both kits require the iPATCORE base system.

Reader Service Number 38

Laser system stores 210,000 pages of data



The MMi-100 optical system from Micro Mart Optisys comes with a Small Computer Systems Interface. It attaches to an IBM PC or compatible using a Host Adapter, which fits into an 8-bit PC slot.

OptiDriver from Micro Mart's Optisys Division allows IBM compatibles to use laser technology to store the equivalent of 1100 floppy disks on a 5.25-inch, reusable cartridge. Users access a WORM (write once, read many) optical disk drive as if it were a Winchester drive. The MMi-100 OptiDriver system can be used with MS-DOS and a variety of word processing programs.

OptiDriver sells for $49.95.

Reader Service Number 37

Color printer interfaces with mainframes

Mitsubishi International's Model CHC-635 nonimpact, thermal-transfer printer features 11-inch horizontal images for its B-size print output. The color printer can produce 270,000 colors for CAD/CAM/CAE hard-copy prints and transparencies.

CHC-635, when interfacing the Shinko Videoprocessor Model SPI 3-1, can be connected with IBM 5080, DEC, and other mainframe color systems. This interface features a five-second capture time. On demand, the video processor can multiplex up to four different terminals.

Contact the company for pricing.

Reader Service Number 39

Reader Interest Survey

Indicate your interest in this department by circling the appropriate number on the Reader Interest Card.

Low 180 Medium 181 High 182

Product Summary

Marlin H. Mickle
University of Pittsburgh

For more information, circle the appropriate Reader Service Number on the Reader Service Card at the back of the magazine.

Manufacturer	Model	Comments	Rs No.
Chips			
Advanced Micro Devices	Supernet chip set	Five-device, bipolar and CMOS chip set implements the ANSI 100M-bit/s FDDI network using fiber-optic media. Replacing six boards of discrete ICs, the LAN can be used for workstations, an interconnection of differing types of networks, and as a link between mainframes, minicomputers, personal computers, and peripherals in a distributed processing environment. $625 in 100-piece quantities.	60
Fujitsu Microelectronics	MB86950 Etherstar controller	LAN controller incorporates protocols for IEEE 802.3 CSMA/CD 10M-bit/s Ethernet and 1M-bit/s Starlan networks. Configurable for 8- or 16-bit bus interfaces, the 1.5-micron, CMOS chip links a host system to a LAN in cost-sensitive network applications connecting personal computers, workstations, disk drives, printers, and other devices. $18.50 in 5000 to 10,000 quantities; 88-pin plastic or ceramic, 84-pin plastic PLCC, or 80-pin flat-pack, surface-mount carrier packaging.	61
Intel Corporation	82385 32-bit controller	CHMOS III cache memory controller supports 20-MHz computers as one of four components in the company's 80386 microprocessor-based Computing Engine chip set. Features of the software-transparent chip include a "posted write" policy that eliminates wait states when writing to main memory and a bus-watching mechanism that ensures cache coherency without performance penalties. $125 each in quantities of 10,000.	62
Boards			
Raster Technologies	GX4000 graphics accelerators	Designed for use with Sun Microsystems workstations, ANSI PHIGS/PHIGS+ board sets plug into a VME backplane to yield 3D graphics for research and scientific 3D visualization, real-time simulation, C3I, geophysics, molecular modeling, and animation.	63
Thomas-Conrad Corporation	TC6046 ARC-Card/MC	Interface board allows users to network IBM PS/2 Micro Channel computers to the Arcnet token-passing LAN. Uses 16K bytes of memory address space, the IBM 16-bit data bus structure, and a dual-port, 2K data buffer. $549 each with two-year warranty.	64
Systems			
Gespac Inc.	Gescomp microcomputers	Real-time, multitasking systems are intended for use as process or cell controllers in factory automation applications, especially with embedded computers in tools and instrumentations. The top-of-the-line 8340-P/HF contains a 32-bit, 16.7-MHz 68020 microprocessor; 68881 arithmetic coprocessor; 2.5M-byte RAM; 1M-byte, 3.5-inch floppy disk; and 40M-byte hard disk. $3995 each, basic configuration; OEM discounts.	65

Manufacturer	Model	Comments	Rs No.
Northwest Instrument Systems	Software Analysis Workstation	National Semiconductor Series 32000 CPU combined with analysis software permits users to picture code behavior during real-time execution in the target system. The SAW system displays information in the symbolic terms of the user's program.	66
Vermont Microsystems	Cobra graphics processor	AutoCad processor draws images at 80,000 clipped vectors/s and offers 1024 × 800 resolution, 16 to 256 simultaneous colors from a 16.7-million palette, and single-screen operation. Optional ADI driver enhances AutoCAD with instantaneous pan and zoom capabilities and full-view inset for parallel real-time viewing. $2995 to $4195, depending on palette selection.	67
Ziltek Corporation	ICE-Engine/68K in-circuit emulator	Provides up to 12.5-MHz, real-time emulation for 68000-series microprocessors; can be connected with IBM PC and compatibles, other host computers, or a dumb terminal through RS-232C ports. Real-time trace memory consists of 63 bits by 4096 words. Up to six levels of nested breakpoints can be used. $5895 with manual and PC driver.	68
Peripherals			
Delkin Devices, USA	525 Extra PS/2 drive	External floppy for Personal System/2 series plugs into existing connector. The 5.25-inch drive allows IBM systems to read, write, and format standard 360K disks as the B drive. $325.	69
Floating Point Systems	P64/210 I/O subsystem	Parallel I/O subsystem for the FPS M64 series controls, collects, and processes data while working in parallel with the CPU. VMS-compatible subsystem can be used to configure the M64 series with multiple units, array processors, DEC computers, high-speed disks, high-density tapes, and high-speed A/D converters. Under $100,000.	70
Software			
Digitalk, Inc.	Smalltalk/V, Version 2.0	PC-based implementation of the Smalltalk object-oriented programming language supports 640 × 480 graphic modes of the IBM PS 2/25 and /30 computers. Bitmapped window enhancements (disk and code browsers, debuggers, free-drawing panes) display windows faster than earlier version. $99.95; upgrades, $25.	71

FREE FACTS. *FAST!*

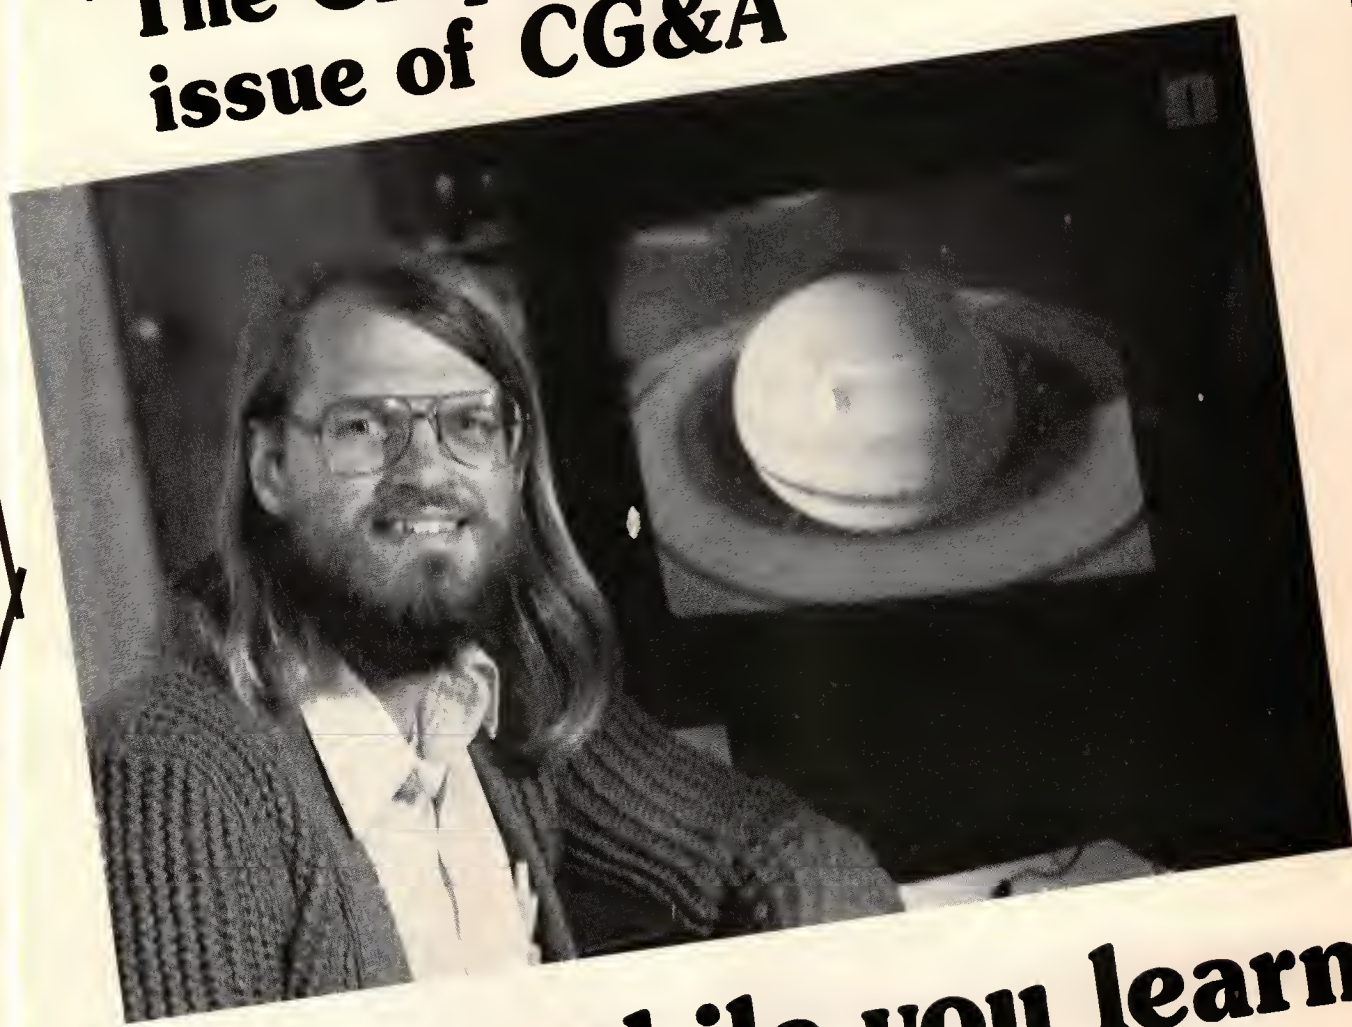
We've just sped up the response time to our Reader Service cards, so now you can get information more quickly about advertisers' products and services and the products we list in our New Products section. Under our new system, the company will have your name and address just days after you send out the card!

While you're indicating the products that interest you, please take a moment to circle the articles and departments that you liked in this issue. That will let us serve your interests better.

We'd like to hear from you!

Jim Blinn's Corner

Don't Miss "The Graphics Man" in every issue of CG&A



Laugh while you learn with the man who created the Mechanical Universe

COMPUTER GRAPHICS & APPLICATIONS
10662 Los Vaqueros Circle, Los Alamitos, CA 90720-9970
Just call (714) 821-8380

Calendar

Conferences sponsored or cosponsored by the Computer Society of the IEEE are indicated by the society's logo. Submit information **eight weeks before cover date** to Calendar, *IEEE Micro,* 10662 Los Vaqueros Circle, Los Alamitos, CA 90720-2578.

February 1988

10th National Computer Conference (IEEE), Feb. 28-Mar. 2, Jeddah, Saudi Arabia. Contact Athar Javaid, IEEE, PO 8900, Jeddah 21492, Saudi Arabia.

Compcon Spring 88, Feb. 29-Mar. 4, San Francisco. Contact Hasan Alkhatib, Dept. of Electrical Engineering and Computer Science, Santa Clara University, Santa Clara, CA 95053; (408) 554-4485.

March 1988

1988 Spring National Design Engineering Show and Conference (materials and components), Mar. 7-10, Chicago, Illinois. Contact Show Manager, Spring National Design Engineering Show, 999 Summer St., Stamford, CT 06905; (203) 964-0000.

Southcon 88 (IEEE, ERA), Mar. 8-10, Orlando, Fla. Contact Southcon 88, 8110 Airport Blvd., Los Angeles, CA 90045; (213) 772-2965.

1988 International Zurich Seminar on Digital Communications, Mar. 8-11, Zurich. Contact Secretariat IZS 88, c/o P. Gunzburger, Hasler AG, TDS, Belpstrasse 23, CH-3000 Bern 14, Switzerland; phone 41 (31) 632-808.

Advances in Semiconductors and Superconductors: Physics and Device Applications, Mar. 13-18, Newport Beach, Calif. Contact SPIE, PO Box 10, Bellingham, WA 98227-0010; (206) 676-3290.

Seventh IEEE Phoenix Conference on Computers and Communications, Mar. 16-18, Scottsdale, Ariz. Contact Carl Ryan, Motorola GEG, 2501 S. Price Rd., Chandler, AZ 85248-2899; (602) 732-3074.

Compstan 88, Computer Standards Conference, Mar. 21-23, Arlington, Va. Contact Roger J. Martin, US Dept. of Commerce, National Bureau of Standards, Technology Bldg. 225, Rm. B266, Gaithersburg, MD 20899; (301) 975-3295.

Sixth IEEE VLSI Test Workshop, Mar. 22-23, Atlantic City, N.J. Contact Wesley E. Radcliffe, IBM East Fishkill, Dept. 277, Bldg. 321-5E1, Hopewell Junction, NY 12533; (914) 894-4346.

April 1988

Computer Networking Symposium, Apr. 11-13, Arlington, Va. Contact George K. Chang, Bell Communications and Research, 6 Corporation Pl., Piscataway, NJ 08854; (201) 699-3879.

CompEuro 88, Apr. 11-15, Brussels. Contact Jacques Tiberghien, VRIJE Universiteit Brussels, Pleinlaan 2, 1050 Brussels, Belgium; phone 32 (02) 641-2905.

Fourth International Conference on Metrology and Properties of Engineering Surfaces, Apr. 13-15, Gaithersburg, Md. Contact K. J. Stout, Coventry Polytechnic, Dept. of Mfg. Systems, Priory St., Coventry CV1 5FB, England; phone 0203 24166, ext. 278; or T.V. Vorburger, A117 Metrology Bldg., National Bureau of Standards, Gaithersburg, MD 20899; (301) 975-3493.

11th IEEE Workshop on Design for Testability, Apr. 19-22, Vail, Col. Contact T.W. Williams, IBM Corp., PO 1900, Dept. 67A/021, Boulder, CO 80301-9191; (303) 924-7692.

Workshop on Microstructure and Macromolecular Research with Cold Neutrons, Apr. 21-22, Gaithersburg, Md. Contact Charles J. Glinka, B106 Reactor Bldg., National Bureau of Standards, Gaithersburg, MD 20899; (301) 975-6242.

May 1988

19th Annual Pittsburgh Conference on Modeling and Simulation, May 5-6, Pittsburgh, Pa. Contact William G. Vogt or Marlin H. Mickle, Modeling and Simulation Conference, 348 Benedum Engineering Hall, University of Pittsburgh, Pittsburgh, PA 15261.

38th Electronic Components Conference (IEEE, EIA), May 9-11, Los Angeles. Contact Electronic Industries Assoc., 2001 Eye St. NW, Washington, DC 20006.

Fifth Workshop on Real-Time Operating Systems, May 12-13, Washington, DC. Contact John A. Stankovic, Dept. of Computer and Information Science, University of Massachusetts, Amherst, MA 01003; (413) 545-0720.

SIGMetrics Conference on Measurement and Modeling of Computer Systems (ACM), May 24-27, Santa Fe, N.M. Contact Connie U. Smith, L and S Computer Technology, 1114 Buckman Rd., Santa Fe, NM 87501; (505) 988-3811.

NCC 88, National Computer Conference (AFIPS, ACM, DPMA, SCS), May 31-June 3, Los Angeles. Contact AFIPS, 1899 Preston White Dr., Reston, VA 22091; (703) 620-8900.

June 1988

ISCAS 88, IEEE International Symposium on Circuits and Systems, June 7-9, Espoo, Finland. Contact Pekka Heinonen, Tampere University of Technology, Computer Systems Laboratory, PO 527, SF-33101 Tampere, Finland, or Olli Simula, Helsinki University of Technology, Dept. of Technical Physics, SF-02150 Espoo 15, Finland.

DAC 88, 25th ACM/IEEE Design Automation Conference, June 12-15, Anaheim, Calif. Contact Pat Pistilli, MP Associates, Inc., 7366 Old Mill Trail, Suite 102, Boulder, CO 80301; (303) 530-4333.

International Conference on Private Switching Systems and Networks (IEE), June 21-23, London. Contact Conference Services Dept., Institution of Electrical Engineers, Savoy Pl., London WC2R OBL, UK

18th International Symposium on Fault-Tolerant Computing, June 27-30, Tokyo. Contact Yasuo Komamiya, 2-4-8 Kikuna, Kohoku-ku, Yokohama 222, Japan; phone (81) 044-911-8181.

July 1988

Navy Micro/OA 88 Conference, July 25-28, San Diego. Contact Code 31.4, NARDAC San Diego, NAS North Island, Bldg. 1482, San Diego, CA 92135-5110; (619) 437-7013.

August 1988

Euromicro 88, 14th Symposium on Microprocessing and Microprogramming, Aug. 29-Sept. 1, Zurich. Contact Chiquita Snippe-Marlisa, PO Box 545, 7500 AM Enschede, The Netherlands; phone 31 (53) 338799.

September 1988

IEEE Artificial Neural Networks Conference, Sept. 18-21, Reston Va. Contact Kamal Karma, 823 Flegler Rd., Gaithersburg, MD 20879; (301) 984-7657.

Advertiser Index

Computer Society Membership Cover IV

Northrop Defense Systems Division 94

Classified Advertising 97

FOR DISPLAY ADVERTISING INFORMATION CONTACT:

Northern California and Pacific Northwest: Don Farris Co., 161 W. 25th Ave., Ste. 102B, San Mateo, CA 94403; (415) 349-2222.
Jack Vance, P.O. Box 3205, Saratoga, CA 95070.
Southern California and Mountain States: Richard C. Faust Co., 24050 Madison St., Ste. 100, Torrance, CA 90505; (213) 373-9604.
Southwest: The House Co., 3000 Weslayan, Ste. 345, Houston, TX 77027; (713) 622-2868.
Midwest: Robert Acker & Associates, 480 Central Ave., Northfield, IL 60093; (312) 441-6050.
East Coast: Hart Associates, PO Box 339, 42 Lake Blvd., Matawan, NJ 07747; (201) 583-8500.
New England: Arpin Associates, PO Box 227, Weston, MA 02193; (617) 899-5613.
Europe: Heinz J. Gorgens, Parkstrasse 8a, D-4054 Nettetal 1 - Hinsbeck (F.R.G); (0 21 53) 8 99 88.

Advertising Director: Dawn Peck.
For production information, conference, and classified advertising, contact Heidi Rex or Marian Tibayan.

IEEE Micro, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720; (714) 821-8380.

Coming in the February issue of *IEEE Micro!*

- The structure and application of videorams
- An overview of R&D efforts in the area of GaAs microprocessors and digital systems design

Moving?

Name (Please Print)

New Address

City

State/Country Zip

☐ Address change: Please notify us 4 weeks in advance

Mail to:
IEEE Micro, Circulation Dept.
10662 Los Vaqueros Cir.
Los Alamitos, CA 90720-2578

ATTACH LABEL HERE

- This address change notice will apply to all IEEE publications to which you subscribe.
- List new address above.
- If you have a question about your subscription, place label here and clip this form to your letter.

Product Index

	RS #	Page#
Boards		
A/D converter	36	99
Controller board	14	94
Coprocessor	24	96
Graphics accelerator	63	100
Interface board	64	100
Plug-in board	20	95
Circuits		
CMOS chip set	60	100
CMOS static RAMs	9	93
Controller	61, 62	100
EPROM product line	5	93
Evaluation module	21	95
In-circuit emulator	68	101
VME engine	25	96
Components		
Digital sound converter	23	96
Data Acquisition		
Data acquisition module	22	96
Organizational system	2	92
Data Communications		
Data communications program	13	94
Digital transmission	1	97
I/O Related Equipment		
Color printer	39	99
Floppy drive	69	101
Image scanner	27	97
I/O subsystem	70	101
Memory/Storage Equipment		
Disk data storage	17	95
Hard-disk controller	28	97
Hard-disk subsystem	34	98
Tape backup	16	95
Publishers		
CS membership	—	Cover IV
Recruitment		
Classified advertising	—	97
Technical professional	—	94
Software		
Application software	19	95
Educational/scientific	10-12	94
Productivity improvement	35	99
Protocol software	15	94
Simulator/debugging	4	92
Smalltalk	71	101
Systems		
Analysis workstation	66	101
Communication network	8	93
Concurrent supercomputers	29	98
Desktop computer	6	93
Graphics processor	67	101
Graphics toolkit	33	98
Hand-held analyzer	26	97
Multiple processor	32	98
Multitasking operating system	31, 65	98, 100
Neurocomputer	18	95
Optical system	37	99
VLSI diagnostic workstation	3	92
Test & Measurement Equipment		
CAE design tools	7	93
Forecasting package	30	98
Software analysis tool	38	99

READER SERVICE CARD

IEEE MICRO

December 1987 issue
(card void after June 1988)

Name_____

Title_____

Company_____

Address_____

City_____ State___ ZIP_____

Country_____ Phone (___) _____

Please send
(Circle those you want):

201 Publications catalog
202 Membership information
203 Student membership information
204 *IEEE Micro* subscription information

Reader interest
(Circle what you liked, add comments on the back)

Readers,
Indicate your interest in articles and departments **by circling the appropriate number** (shown on the last page of articles/departments) **in the shaded section of this card under Product Inquiries.**

Product inquiries **1**
(circle the numbers for products and advertisers you want more information on)

1	21	41	61	81	101	121	141	161	181
2	22	42	62	82	102	122	142	162	182
3	23	43	63	83	103	123	143	163	183
4	24	44	64	84	104	124	144	164	184
5	25	45	65	85	105	125	145	165	185
6	26	46	66	86	106	126	146	166	186
7	27	47	67	87	107	127	147	167	187
8	28	48	68	88	108	128	148	168	188
9	29	49	69	89	109	129	149	169	189
10	30	50	70	90	110	130	150	170	190
11	31	51	71	91	111	131	151	171	191
12	32	52	72	92	112	132	152	172	192
13	33	53	73	93	113	133	153	173	193
14	34	54	74	94	114	134	154	174	194
15	35	55	75	95	115	135	155	175	195
16	36	56	76	96	116	136	156	176	196
17	37	57	77	97	117	137	157	177	197
18	38	58	78	98	118	138	158	178	198
19	39	59	79	99	119	139	159	179	199
20	40	60	80	100	120	140	160	180	200

READER SERVICE CARD

IEEE MICRO

December 1987 issue
(card void after June 1988)

Name_____

Title_____

Company_____

Address_____

City_____ State___ ZIP_____

Country_____ Phone (___) _____

Please send
(Circle those you want):

201 Publications catalog
202 Membership information
203 Student membership information
204 *IEEE Micro* subscription information

Reader interest
(Circle what you liked, add comments on the back)

Readers,
Indicate your interest in articles and departments **by circling the appropriate number** (shown on the last page of articles/departments) **in the shaded section of this card under Product Inquiries.**

Product inquiries **2**
(circle the numbers for products and advertisers you want more information on)

1	21	41	61	81	101	121	141	161	181
2	22	42	62	82	102	122	142	162	182
3	23	43	63	83	103	123	143	163	183
4	24	44	64	84	104	124	144	164	184
5	25	45	65	85	105	125	145	165	185
6	26	46	66	86	106	126	146	166	186
7	27	47	67	87	107	127	147	167	187
8	28	48	68	88	108	128	148	168	188
9	29	49	69	89	109	129	149	169	189
10	30	50	70	90	110	130	150	170	190
11	31	51	71	91	111	131	151	171	191
12	32	52	72	92	112	132	152	172	192
13	33	53	73	93	113	133	153	173	193
14	34	54	74	94	114	134	154	174	194
15	35	55	75	95	115	135	155	175	195
16	36	56	76	96	116	136	156	176	196
17	37	57	77	97	117	137	157	177	197
18	38	58	78	98	118	138	158	178	198
19	39	59	79	99	119	139	159	179	199
20	40	60	80	100	120	140	160	180	200

SUBSCRIPTION CARD

SUBSCRIBE TO IEEE MICRO

☐ YES, sign me up! ☐ Renew my subscription!

If you are a member of the Computer Society or any other IEEE society, pay the member rate of only **$17** for a year's subscription (6 issues). Subscriptions are annualized with membership. For orders submitted March through August, please pay half the given full-year fee for a half-year's subscription.

Society:_____ IEEE membership no:_____

☐ YES, sign me up! ☐ Renew my subscription!

If you are a member of ACM, NSPE, ACS, IEE (UK), SCS, IPSJ, IECEJ, or any other technical society, pay the sister society rate of only **$25** for a year's subscription (6 issues).

Society:_____ Mem. no. (if any):_____

Full Signature *Date*

Name

Street

City

State/Country *ZIP/Postal Code*

☐ Payment enclosed .

☐ Charge to ☐ Visa ☐ MasterCard ☐ AmEx

☐☐☐☐☐☐☐☐☐☐☐☐☐☐☐☐

Charge Card Number

Mo. Yr.
☐☐ ☐☐
Exp. Date

M1287

Charge orders also taken by phone:
(714) 821-8380 8:00 a.m. to 5:00 p.m. Pacific time
Circulation Dept.
10662 Los Vaqueros Cir.
Los Alamitos, CA 90720

Editorial comments

I liked: _____

I disliked: _____

I would like to see: _____

For reader service inquiries, use other side

PO box address for reader service cards only

PLACE STAMP HERE

IEEE MICRO

Reader Service Inquiries
PO Box 16508
North Hollywood, CA 91615-6508
USA

Editorial comments

I liked: _____

I disliked: _____

I would like to see: _____

For reader service inquiries, use other side

PO box address for reader service cards only

PLACE STAMP HERE

IEEE MICRO

Reader Service Inquiries
PO Box 16508
North Hollywood, CA 91615-6508
USA

NO POSTAGE NECESSARY IF MAILED IN THE UNITED STATES

BUSINESS REPLY MAIL
FIRST CLASS PERMIT NO. 38 LOS ALAMITOS, CA

POSTAGE WILL BE PAID BY ADDRESSEE

Computer Society of the IEEE
Circulation Dept.
10662 Los Vaqueros Cir.
Los Alamitos, CA 90720-9970

THE COMPUTER SOCIETY

A member society of the Institute of Electrical and Electronics Engineers, Inc.

Executive Committee

President: Roy L. Russo*
IBM T.J. Watson Research Center
PO Box 218, Route 134
Yorktown Heights, NY 10598
(914) 945-3085

President-Elect: Edward A. Parrish, Jr.*

Vice Presidents
Education: Michael C. Mulder (1st VP)*
Technical Activities: Kenneth R. Anderson (2nd VP)*
Area Activities: Willis K. King†
Conferences and Tutorials: James H. Aylor†
Membership and Information: Merlin G. Smith†
Publications: J.T. Cain†
Standards: Helen M. Wood

Secretary: Duncan H. Lawrie
Treasurer: Joseph E. Urban
Past President: Martha Sloan*
Division V Director: Martha Sloan
Division VIII Director: H. Troy Nagle†
Executive Director: T. Michael Elliott†

*voting member of the Board of Governors
†nonvoting member of the Board of Governors

Board of Governors

Term Ending 1987	Term Ending 1988
Barry W. Boehm	Mario R. Barbacci
Paul L. Borrill	Victor R. Basili
Glen G. Langdon, Jr.	Lorraine M. Duvall
Duncan H. Lawrie	Michael Evangelist
Susan L. Rosenbaum	Allen L. Hankinson
Bruce D. Shriver	Laurel Kaleda
Harold S. Stone	Ted Lewis
Wing N. Toy	Ming T. Liu
Helen M. Wood	Earl E. Swartzlander, Jr.
Akihiko Yamada	Joseph E. Urban
Oscar N. Garcia†	

Next Board Meeting

8:30 a.m.-5 p.m., March 4, 1988
Cathedral Hill Hotel, San Francisco

Senior Staff

Executive Director: T. Michael Elliott
Editor and Publisher: True Seaborn
Director, Computer Society Press: Chip G. Stockton
Director, Conferences: William R. Habingreither
Director, Finance and Administration: Mary Ellen Curto

Offices of the Computer Society

Headquarters Office
1730 Massachusetts Ave. NW
Washington, DC 20036-1903
General Information: (202) 371-0101
Publications Orders: (800) 272-6657
Telex: 7108250437 IEEE COMPSO

Publications Office
10662 Los Vaqueros Circle
Los Alamitos, CA 90720
Membership and General Information: (714) 821-8380

European Office
13, Avenue de l'Aquilon
B-1200 Brussels, Belgium
Phone: 32 (2) 770-21-98
Telex: 25387 AVVALB

Purpose

The Computer Society strives to advance the theory and practice of computer science and engineering. It promotes the exchange of technical information among its 90,000 members around the world, and provides a wide range of services which are available to both members and nonmembers.

Membership

Members receive the highly acclaimed monthly magazine *Computer*, discounts on all society publications, discounts to attend conferences, and opportunities to serve in various capacities. Membership is open to members, associate members, and student members of the IEEE, and to non-IEEE members who qualify as affiliate members of the Computer Society.

Publications

Periodicals. The society publishes six magazines (*Computer, IEEE Computer Graphics and Applications, IEEE Design & Test of Computers, IEEE Expert, IEEE Micro, IEEE Software*) and three research publications (*IEEE Transactions on Computers, IEEE Transactions on Pattern Analysis and Machine Intelligence, IEEE Transactions on Software Engineering*).

Conference Proceedings, Tutorial Texts, Standards Documents. The society publishes more than 100 new titles every year.

Computer. Received by all society members, *Computer* is an authoritative, easy-to-read monthly magazine containing tutorial, survey, and in-depth technical articles across the breadth of the computer field. Departments contain general and Computer Society news, conference coverage and calendar, interviews, new product and book reviews, etc.

All publications are available to members, nonmembers, libraries, and organizations.

Activities

Chapters. Over 100 regular and over 100 student chapters around the world provide the opportunity to interact with local colleagues, hear experts discuss technical issues, and serve the local professional community.

Technical Committees. Over 30 TCs provide the opportunity to interact with peers in technical specialty areas, receive newsletters, conduct conferences, tutorials, etc.

Standards Working Groups. Draft standards are written by over 60 SWGs in all areas of computer technology; after approval via vote, they become IEEE standards used throughout the industrial world.

Conferences/Educational Activities. The society holds about 100 conferences each year around the world and sponsors many educational activities, including computing sciences accreditation.

European Office

This office processes Computer Society membership applications and handles publication orders. Payments are accepted by cheques in Belgian francs, British pounds sterling, German marks, Swiss francs, or US dollars, or by American Express, Eurocard, MasterCard, or Visa credit cards.

Ombudsman

Members experiencing problems — late magazines, membership status problems, no answer to complaints — may write to the ombudsman at the Publications Office.

Information

Use the Reader Service Card to obtain the following material:

- Membership information and application (RS #202)
- Publications catalog (proceedings, tutorials, standards) (RS #201)
- Periodicals subscription application/information for individuals (members, sister-society members, others) (RS #200)
- Periodicals subscription application/information for organizations (libraries, companies, etc.) (RS #199)
- List of awards and award nomination forms (RS #198)
- Technical committee list and membership application (RS #197)
- Directory of officers, board members, committee chairs, representatives, staff, chapters, standards working groups, etc. (RS #196)

JOIN US.

THE COMPUTER SOCIETY OF THE IEEE

THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.

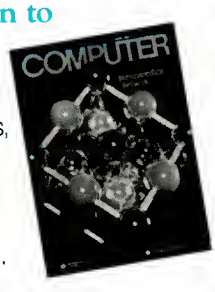


The Computer Society gives you the balance you need.

From books on artificial intelligence to conferences on design automation, the Computer Society has the information you want and the balance between theory and practice you need.

We invite you to join us and receive these benefits:

An automatic subscription to COMPUTER Magazine



This authoritative yet readable monthly brings you industry surveys, tutorials, and application perspectives across the breadth of the computer field—from theory to practice, and every step in between.

Low member rates on other Computer Society magazines and transactions

Our magazines provide a practical application-oriented treatment of the leading edge of computer technology and our transactions provide the more analytical and theoretical base—the perfect balance between theory and practice.



The opportunity to work shoulder-to-shoulder with experts

in one or more of our standards working groups (we've got more than 80 now!) to develop draft standards, which after approval, become IEEE standards used throughout the world. (The Computer Society is unique in offering its members the opportunity to develop computer-related standards.)

Participate in one or more of our 33 technical committees—

networks of professionals with common interests in computer hardware, software, and applications. Technical committee programs are determined by the membership. Some are more theoretical; others are more applied.

Discounts on conference registration fees

Choose from more than 100 worldwide annually, from large applications-oriented conferences to small workshops, and a range of theory and practice for every interest.

To obtain membership information and an application, circle Reader Service Card #202.